AD-A039 548 GTE SYLVANIA INC NEEDHAM HEIGHTS MASS ELECTRONIC SYS--ETC F/G 17/2
SENET-DAX STUDY. VOLUME 1.(U)
JUN 76 DCA100-75-C-0071

UNCLASSIFIED FR76-1-VOL-1 NL

1 OF 3
AD
A039548

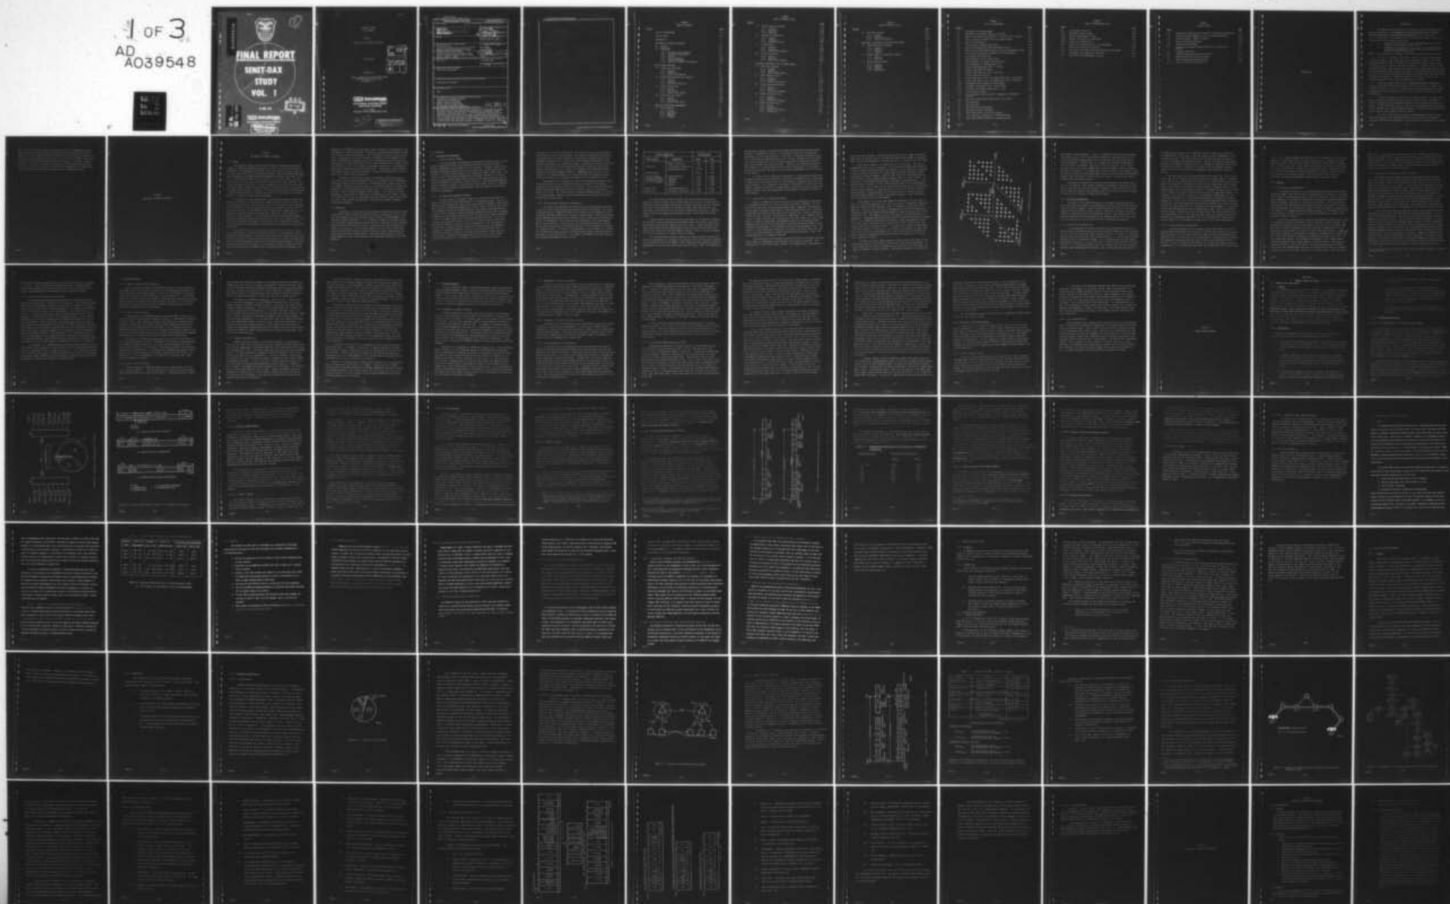

1

B.S.

FINAL REPORT

SENET-DAX STUDY VOL. 1

25 JUNE 1976

GTE SYLVANIA
INCORPORATED
ELECTRONIC SYSTEMS GROUP
EASTERN DIVISION

SENET-DAX STUDY

FINAL REPORT

Volume 1

(Contract No. DCA-100-75-C-0071)

25 June 1976

Submitted to

Defense Communications Engineering Center
Defense Communications Agency
Washington, D.C.

GTE SYLVANIA
INCORPORATED
ELECTRONIC SYSTEMS GROUP
EASTERN DIVISION

77 "A" STREET
NEEDHAM HEIGHTS, MASSACHUSETTS 02194

See 1473

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER

2. GOVT ACCESSION NO.

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle)

6 ~~FINAL REPORT~~
SENET-DAX STUDY. Volume 1.
~~VOL. 1 AND VOL. 2~~

5. TYPE OF REPORT & PERIOD COVERED

9 Final Report.
June 75 — June 76.

6. PERFORMING ORG. REPORT NUMBER

14 FR76-1-Vol-1

7. AUTHOR(s)

8. CONTRACT OR GRANT NUMBER(s)

15 DCA 100-75-C-0071

9. PERFORMING ORGANIZATION NAME AND ADDRESS

GTE Sylvania Inc.
Electronic Systems Group, Eastern Division
Needham Heights, Massachusetts 02194

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS

PE 33143K
Task 13103G

11. CONTROLLING OFFICE NAME AND ADDRESS

Advanced Systems Concepts Branch, R740
Defense Communications Engineering Center
Reston, Virginia 22090

12. REPORT DATE

11 25 June 1976

12 257p.

13. NUMBER OF PAGES

490

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)

Same as 11

15. SECURITY CLASS. (of this report)

Unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

Distribution unlimited.

18. SUPPLEMENTARY NOTES

None.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Integrated voice/data switching
Time division multiplexing
Dynamic channel allocation
Digital Access Exchange (DAX)
Distributed Processor Architecture

406 451 B

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The SENET-DAX study investigated the concept of an integrated voice/data switching system using time division multiplexing and dynamic channel allocation; the objectives being: (1) To analyze the feasibility of the SENET-DAX concept in terms of structure and performance characteristic. (2) To address the translation of the concept into candidate hardware and software techniques that could be used to implement the concept. The results of the analysis and techniques studies is presented in these two volumes.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

VOLUME 1

TABLE OF CONTENTS

VOLUME 1

TABLE OF CONTENTS (Cont)

VOLUME 1

TABLE OF CONTENTS (Cont)

VOLUME 1

LIST OF ILLUSTRATIONS

VOLUME 1

LIST OF ILLUSTRATIONS (Cont)

VOLUME 1

LIST OF TABLES

INTRODUCTION

INTRODUCTION

The SENET-DAX Study described in this report is directed towards investigation of the concept of an integrated voice/data switching system using time division multiplexing and dynamic channel allocation as first suggested by Drs. G. Coviello and P. Vena of DCEC in [19], and further developed by them in [20]. The objectives of this study task have been:

a. To analyze the feasibility of the SENET-DAX concept in terms of structure and performance characteristics

b. To address the translation of the concept into candidate hardware and software techniques that could be used for the conceptual design of an access area exchange

Results of the analysis and techniques studies are presented in these two volumes.

Section 1 is a summary section that considers the impact of the dynamic allocation concept on those telecommunications areas where the most interesting results have come to light: the precedence and preemption scheme, speech security, call routing, service features, and network management. Sections 2 and 3 describe in detail the master frame structure for digitized voice, packet data, and bulk data, and the requirements for processing and control of the integrated voice/data stream. An analysis of systems requirements for signaling, synchronization, and error control is presented in Section 4. Functional allocations, covered in Section 5, are initially examined for their impact on SENET-DAX architecture.

Tradeoffs towards selection of the distributed processor architecture, considered to be the most promising approach for realization of the concept, are covered in Section 6. Sections 7, 8 and 9 contain detailed descriptions of software and processor structure and synchronization and error control techniques for this architecture.

Section 10 presents the results of performance analysis, including flexibility for interfacing with dissimilar trunk and terminal devices, the amount of control overhead requiring transmission capacity, and an analysis of blocking and delays. Section 11 provides a general assessment of system capabilities with regard to size and modularity, traffic handling capability, availability, interoperability, speech security, and system and technical control.

One year is insufficient to study to the maximum every aspect of performance and technique for a new and comprehensive concept such as the SENET-DAX. Study areas that have been identified as critical for basic evaluation of the concept,

such as control signaling, software and processor structure, synchronization technique, and delay and blocking analysis, have been studied and described in detail. Other areas, such as COMSEC processing (including key distribution) and digital voice conversion, have been addressed but remain to be studied in more depth. Areas such as satellite link performance require initial investigation. In addition, there is need for a verification in software and hardware of the techniques that have been postulated. Section 12 concludes the report with recommendations for further analysis and development towards eventual realization of the SENET-DAX concept.

SECTION 1

THE IMPACT OF DYNAMIC ALLOCATION

SECTION 1

THE IMPACT OF DYNAMIC ALLOCATION

1.1 PROBLEM

The SENET-DAX concept differs from conventional transmission multiplexing schemes in allocating transmission capacity to a call in accord with the specific information rate requirements of the call, and in allowing the location of the capacity allocated to the call to be dynamically variable. Both conventional SDMX (Space Division Multiplex) and TDMX (Time Division Multiplex) transmission systems, by contrast, allocate fixed portions of their capacity (a circuit, for SDMX; a specific allocation of bit slots, equal to all other allocations, for TDMX) regardless of the actual requirements of the call. In fact, for TDMX, the bit rate of the subscriber instruments must be matched to that fixed by the transmission system, often resulting in further inefficiencies. In neither of the traditional broad types of multiplexing is there found any usefulness in dynamic variation of the location of the allotment made to a call after the call has been connected.

A less obvious common characteristic of conventional multiplex transmission systems in current use is the necessity to commit transmission system resources to an end-to-end connection. Prior to subscriber use of the connection, this commitment is fully adequate to serve the subscribers involved, regardless of whether the connection will be used. In the SDMX case this is necessary for inband signaling, even though such signaling ordinarily requires only a small part of the transmission capability of the connection. In a conventional TDMX system using Common Channel Interswitch Signaling, this would not be necessary, but is still commonly done. Moreover, the CCIS channel is frequently one of the full-sized communication channels, reserved full-period for this purpose, and unavailable, even when idle, for subscriber connections.

Conventional line or circuit-switching systems now in use for voice transmission and switching in military and other governmental systems frequently provide for preemption by precedence. In general, these systems have handled facsimile transmissions as pseudo-voice connections, and any precedence involved was associated with the level specified during the initial voice-instrument signaling process. Similarly, while narrative/record data traffic has been message-switched in conventional Store-and Forward systems such as AUTODIN, with an established scheme for precedence and

preemption, the AUTODIN was not designed to handle interactive, query/response, data base update, or bulk data transfers. Consequently, no universal precedence structure for all types of data transmissions has been established, nor were satisfactory protocols established for the variety of modes of operation to be accommodated all at once. The ARPANET was designed to improve this transmission involving interactive, query/response and data base updates, and has been demonstrated to be capable of bulk data transfers, but utilizes a simplistic precedence scheme, and is not designed to operate with conventional narrative/record traffic. Additionally, neither AUTODIN nor ARPANET provides a complete treatment of message security appropriate to all of these various types of data transmission.

In addition to the potential inpact of the dynamic allocation concept on transmission capacity and the precedence scheme, the integration of non-homogeneous subscriber traffic and its transmission and switching may well impact other areas. For examples, routing of voice and pseudo-voice connections might use a common technique, but it will not necessarily be optimal for all (or any) of the data transmissions; service feature requirements differ widely; with respect to network management, a simple count of connections in conventional systems equates to usage, but in SENET-DAX, it does not, particularly when the system is simultaneously handling voice, video, and data transmissions; and subscriber terminal requirements could well be affected.

1.2 OBJECTIVES

The objectives of this section are to discuss and scope the technical impact of the dynamic channel allocation concept on various system aspects such as precedence and preemption, security, routing, service features, network management, and subscriber terminal requirements. It is intended to summarize the insights gained into the dynamic channel allocation concept during the study period and to highlight significant areas of interest and concern. Specific details can be found in the balance of this report. In this section, emphasis will be placed on differences between the SENET-DAX concept and the more traditional types of multiplexing concepts, and particularly on those aspects in which the dynamic channel allocation concept appears to offer clear advantages over conventional approaches to the same problems.

1.3 DISCUSSION

1.3.1 Precedence and Preemption

1.3.1.1 Current Precedence Schemes

A system of precedence designations and preemption protocols is now in use in military analog and digitized voice transmission and switching systems such as AUTOVON, AUTOSEVOCOM, and the AN/TTC-39 Circuit Switches. Presently, however, these precedence/preemption usages do not cover low-speed video, forward-error-corrected facsimile, or sensor bulk data being circuit-switched via SDMX or TDMX systems. Similarly, a system of precedence designations and preemption protocols is in use in military analog-equivalent digital data transmission and switching systems such as AUTODIN and ARPANET (the latter not strictly military). These precedence/preemption usages are primarily limited to narrative/record data traffic, and only a simplistic approach exists for interactive, query/response, data base update, and non-sensor bulk data transmissions.

1.3.1.2 Precedence and Preemption Overview

The SENET concept is intended to handle all types of voice, pseudo-voice, and data transmissions, often simultaneously, in an integrated manner. Thus it becomes necessary to extend the virtual real time precedence structure to include both digitized voice and bulk data transferred in Class I, and to extend the data switching precedence structure to include interactive, data update, and other bulk data transfers, and to reconcile the two resulting precedence structures with each other. An additional consideration from this study is the necessity to define a precedence structure for Common Channel Interswitch Signaling (CCIS) messages. These CCIS messages, which concern both Class I and Class II transmissions, have been found best carried through the system as Class II transmissions. The CCIS precedence structure must be reconciled with the previously identified pair of precedence structures. Evaluation and analysis has led to the conclusions that types of traffic not previously subject to a precedence and preemption scheme should use the precedence designations already developed for the class of traffic within which they naturally fall. For example, forward-error corrected facsimile

transmission over a virtual connection in real time would use a voice precedence level during the initial signaling phase. By contrast, an ARQ-protected facsimile transmission sent via packet-switched store-and-forward techniques would use the appropriate precedence level of that type of data transmission (generated at the subscriber-instrument) in a transmission header. Precedences of virtual connections would be compared to each other during the queueing of packets for transmission. However, high precedence (e.g., CRITIC, ECP, or FLASH) data transmissions should preempt low precedence (ROUTINE, PRIORITY, or IMMEDIATE) voice or other virtual connections even though preemption by a relatively brief data transmission could result in termination of a typically much lengthier voice or pseudo-voice call.

With respect to CCIS messages, we feel that each should be assigned a data message precedence equal to the precedence of the call that the CCIS message controls, with the exception of Class I call release. After a call has terminated, better SENET network performance results if channel resources are reclaimed as quickly as possible. This could be done by assigning a high precedence to each CCIS call release message. However, alternative methods may be available that minimize preemption of on-going calls.

1.3.1.3 Data Precedence for Class II Transmission

Our analysis has indicated that extension of the voice precedence structure to cover non-voice virtually-connected bulk data would be satisfactory. With regard to data procedures, it appeared initially that extension of the AUTODIN-type narrative/record precedence structure to interactive, query/response, and data base update packet-switched transmissions posed a problem because of their greatly differing response time requirements from narrative/record transmission such as ARQ-facsimile and non-sensor bulk data. Delivery delays for fast-response interactive, query/response, and data base update are normally required to be no more than a few seconds to a couple of minutes. By contrast, delivery delays of a few minutes to a few hours can be permitted for narrative/record and ARQ-facsimile traffic, and many minutes to many hours for bulk data transfers. For example, the AUTODIN-II specification shows the following requirements for that system.

TYPE OF TRANSACTION		DELIVERY DELAY	
MAIN CATEGORY	SUBCATEGORY	MEAN	MAX.
Interactive	Human Interaction	1 SEC	2 SEC
	Alarms/Status Indicators	1 SEC	3 SEC
	Monitoring/Telemetry	0.3 SEC	1 SEC
Query/Response		36 SEC	2 MIN
Narrative/Record (Including AUTODIN-I Transmission Times of 2 to 20 minutes)	CRITIC/ECP/FLASH	4 MIN	10 MIN
	IMMEDIATE	6 MIN	20 MIN
	PRIORITY	9 MIN	3 HRS
	ROUTINE	16 MIN	6 HRS
Bulk-1 Data	Files/Programs/Results	5 MIN	20 MIN
Bulk-2 Data	Data Bases, Sensor Data	4 HRS	12 HRS

It seemed initially that if delivery time criteria for the fast-response transmissions were to be met, these criteria would be excessively stringent for the longer-delay narrative/record and bulk data transfers, resulting in an over-engineered system. On the other hand, if criteria for longer-delay items were met, even comfortably, there would be little likelihood of meeting the interactive, query/response, and data base update criteria except at times of very light data traffic loading.

Further analysis has indicated that this problem can be resolved by handling the non-Class I data transmission in two different modes of operation: interactive, query/response, and data base update traffic will utilize a "store-and-squirt" technique; and narrative/record, ARQ-facsimile, and non-sensor bulk data traffic will utilize a more classical "store-and-forward" technique. Details of these techniques will be found in Section 3.2, Traffic Control Procedures.

The store-and-squirt technique requires only momentary packet accountability during the transmission of a message to an addressee who has been previously established as available and able to be reached. Thus, transmission time through the SENET is approximately the sum of cross-office delays and link transit times, overlapped (in a

multiple-packet message or response) with message receipt by the originating SENET-DAX from the originating subscriber, and by message delivery by the terminating SENET-DAX to the receiving subscriber. By contrast, in the store-and-forward technique, the originating SENET-DAX must receive, validate, and make itself accountable on a complete message basis for the entire message from an originating subscriber prior to any transmission through the SENET. Similarly, the SENET-DAX on which the first addressee is terminated must receive, validate, and make itself accountable for the entire message, and transmit an end-to-end acknowledgement prior to delivery of the message.

As a result of this difference in treatments, it is possible that those data messages utilizing the store-and-forward transmission technique will have significantly longer cross-network delays than will those messages transmitted in the store-and-squirt technique. As a consequence, there appears no necessity to conceive of a precedence structure for interactive, query/response, and data base update messages that would be different from (and ranked above) the precedence structure for narrative/record traffic.

1.3.1.4 Reconciliation of Precedence Schemes

Figure 1-1 summarizes the current status of the recommendations offered for the interrelationship of voice, data, and control precedence structures. The left half of the figure shows the interrelationships found or suggested for precedence of data packets requiring transmission in a Store-and-Squirt mode or in a Store-and-Forward mode with respect to each other and to existing Class I virtual connections. The qualification "existing" is intended to be significant. It is suggested that if a Class I virtual connection must be allocated at the same instant as a Class II data packet is ready for transmission, and SENET-DAX capacity for the upcoming frame is adequate for only one of the two requirements, then the data packet should be given precedence over the Class I call, regardless of the relative precedence levels. The result is a single frame delay added to the cross-office connection time of the Class I call, while the data packet is removed from the transmission queue. This is considered preferable to the opposite approach.

The direction of the arrow shown at each intersection in the figure indicates which traffic is recommended to take precedence if a choice is necessary. Thus, existing Class I calls should always continue if the queued Class II data packet is

of equal or lower precedence, but CRITIC data packets are of an order of precedence higher than any used for Class I calls and hence would always preempt or interrupt Class I calls of any precedence if transmission space is necessary in the next frame.

Note that the figure is not complete for all intersections. The areas of incompleteness fall into two categories. For the plane representing coincidence of Store-and-Forward and Store-and-Squirt data packets, one area represents coincidence of equal precedence data packets. It seems likely that Store-and-Forward traffic should defer to Store-and-Squirt traffic of equal precedence, considering the differences found in their probable required response times (or delivery delays). The second area is one diagonal removed from the diagonal of equal precedence. If the same precedence structure is to cover both Store-and-Squirt and Store-and-Forward Class II traffic, the decisions in this second category would seem to be both obvious and unavoidable. However, it seems likely that a case can be made for permitting Store-and-Squirt of given precedence to receive preference in transmission over Store-and-Forward traffic of even one level higher precedence, again considering their relative delivery delay requirements.

With respect to the planes that represent coincidence of the two types of Class II traffic with on-going Class I calls, the areas of incompleteness shown can again be divided into two categories, similar to although not identical with the above two. In these cases, the areas of incompleteness include the diagonals once-removed and those twice-removed from the diagonal of equal precedence, both towards the Class II axis. The argument in these cases hinges on the differing effect of preemption. A preempting data packet may terminate (or merely interrupt) an on-going Class I call. Affording preferential continuation of that lower precedence Class I call simply delays the transmission of the Class II traffic. Within reason, and quite probably with different tolerances with respect to the nature and precedence of the data packets, or to that of the call likely to be preempted, it seems likely that better overall system performance will result if data traffic preemption of Class I traffic is accomplished judiciously rather than ruthlessly as in current voice preemption techniques.

The right side of the figure represents similar results with respect to the coincidence of Class II data traffic (of both major types) with CCIS messages. It will be noted that the areas of incompleteness on the right side of the figure are simply the diagonals of equal precedence. Contradictory arguments can be advanced

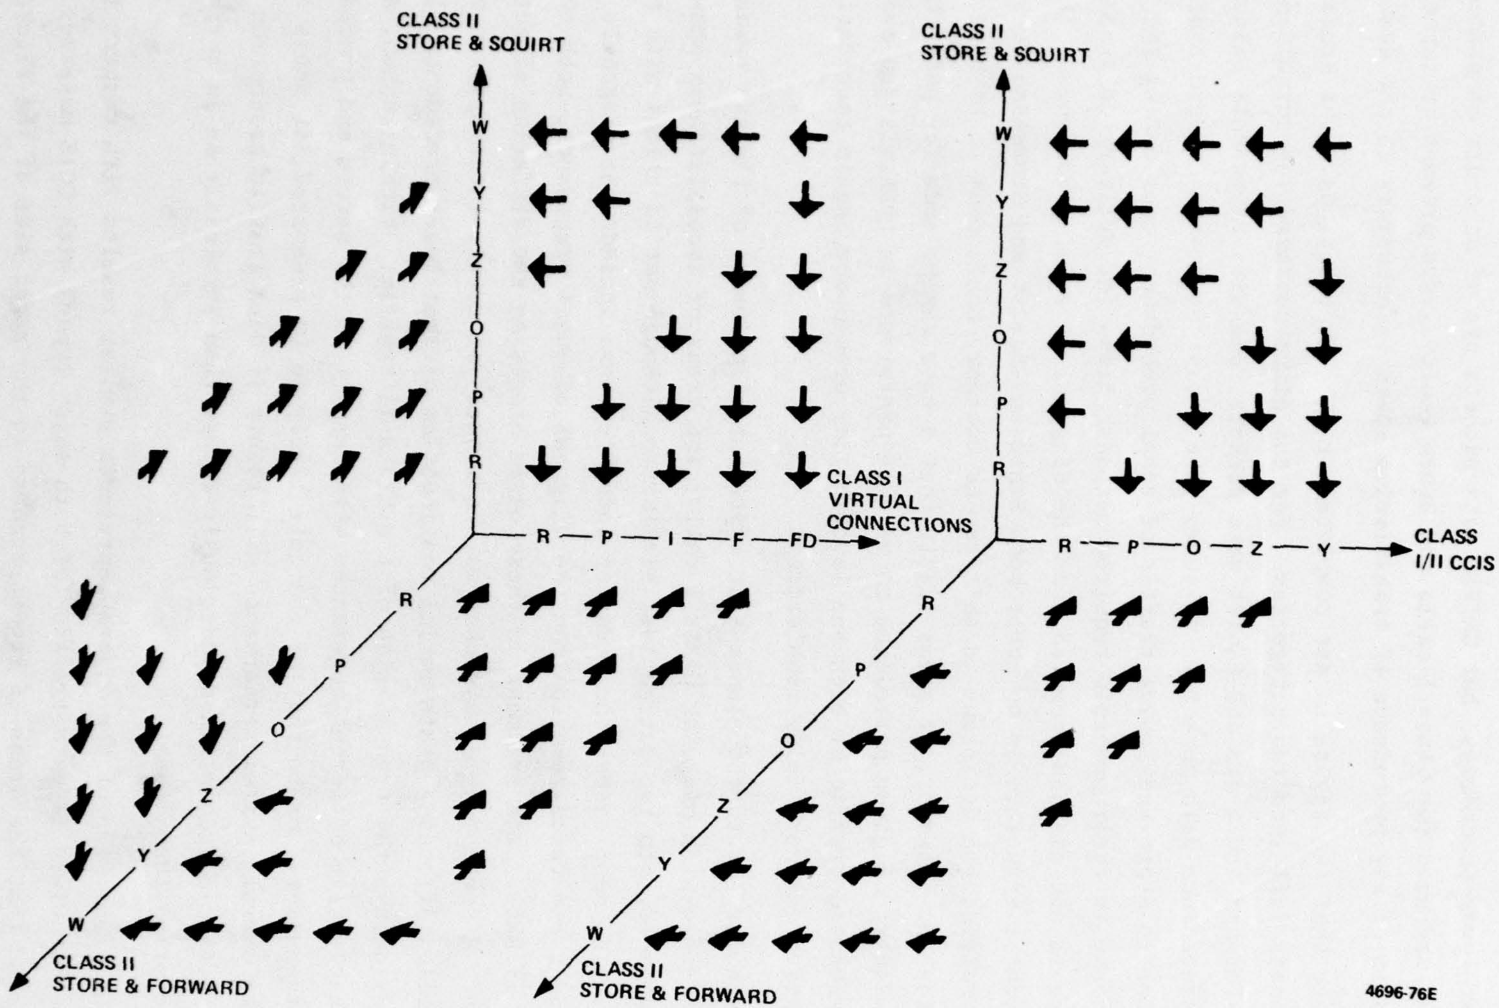

Figure 1-1. Precedence Inter-Relationships

with respect to these diagonals. First, it is arguable that both Store-and-Forward and Store-and-Squirt Class II data traffic represent the subscriber's use of the system and that these usages should be given preference over the system's own internal requirements. Alternatively, it is arguable that efficient maintenance of the CCIS traffic is preferable, not only to service any one call, but to service efficiently all calls in this "common user" system environment. Finally, and perhaps most convincingly, it is arguable that the Store-and-Squirt/CCIS diagonal should be resolved in favor of the Store-and-Squirt traffic, while the Store-and-Forward/CCIS diagonal should be resolved in favor of the CCIS traffic, considering the relative delivery delay requirements involved.

Resolution of incomplete intersections in both halves of the figure remains a task to be addressed. The ultimate precedence and preemption scheme will be primarily a function of user desires, and only secondarily a function of SENET-DAX capability, since the powerful stored-program control in the system architecture appears to lend itself to a variety of reconciliation schemes.

1.3.2 Security Considerations

The impacts of the SENET-DAX concept on Call/Crypto synchronization and on end-to-end delays attributable to crypto signaling and control have not yet been completely evaluated. However, initial consideration given to these areas has already indicated that crypto effects appear likely to affect two of the more interesting capabilities made possible due to the SENET-DAX concept. These capabilities, both directly attributable to the flexibility of the SENET-DAX concept, are the possibilities of digitized voice conversion and of Time-Assigned-Data-Interpolation.

1.3.2.1 Digitized Voice Conversion

It is assumed that the likeliest application of the SENET-DAX concept will result in the system handling both secure and nonsecure transmissions, quite possibly in both classes of traffic, but almost certainly in Class I. Moreover, it is assumed that initially most secure calls will be Crypto Stage I, i.e., link-by-link encrypted, probably in bulk over interswitch trunks. As the system and the subscriber equipments mature and evolve, it is likely that more and more transmissions will be Crypto Stage II, i.e., end-to-end encrypted, after link-encryption or clear text signaling and bulk super-encryption over interswitch trunks. Nothing has been noted so far in

the evaluation and analysis of the SENET-DAX concept to indicate any difficulties in handling clear, Crypto Stage I, and Crypto Stage II calls, independently or simultaneously or both. This is particularly true of digitized voice calls using a variety of techniques and transmission speeds. However, the potential simultaneous existence of incompatible digitized voice calls through the SENET-DAX network raises the question of whether a capability to translate from one technique or speed to another, similar to the capability for incompatible data subscribers, should be offered.

The translation issue is discussed further in Section 1.3.4. If it is feasible to map directly from one digitization algorithm to another (which is a separate issue), and if such a feature were to be offered, it is clear from the security standpoint that translation is feasible for both clear text transmissions and for Crypto Stage I calls, which are clear text within the switches traversed by the call. It is equally apparent that voice conversion could not be offered for Crypto Stage II calls, since by definition, switches would be unable to detect the appropriate symbols to be converted. It is possible to conceive of an intermediate Crypto Stage I/II call, in which the transmission is end-to-end encrypted from the originating subscriber equipment to a single (e.g., originating or terminating) SENET-DAX switch, decrypted at that switch only long enough for voice conversion, and then end-to-end re-encrypted from the converting switch to the terminating subscriber equipment. If voice conversion algorithm mapping were made a trusted process under a multi-level-security software scheme, vulnerability of the call would seem to be not much worse than for a true Crypto Stage II call. Note that this Crypto Stage I/II treatment would only be necessary when incompatible subscribers were to be interconnected; compatible subscribers could still be connected as Crypto Stage II calls as the system matures.

1.3.2.2 Time-Assigned Data Interpolation

A situation analogous to voice conversion exists with respect to the Time-Assigned-Data-Interpolation (TADI) capability. Briefly, this concept (discussed more fully in Section 1.3.5) involves detection of unused directions of transmission in digitized voice calls and of pauses in on-going speech. These portions of allocated but temporarily unused capacity could then be used for the transmission of data packets. If such a feature were to be offered, it is again clear from the security standpoint that it is feasible for both clear text transmissions and for Crypto Stage

I calls. It is equally apparent that TADI could not be accomplished in either Crypto Stage I/II (as described above) or in Crypto Stage II calls. Consequently, if the SENET-DAX system is likely to evolve steadily toward full Crytpo Stage II handling for all digitized voice calls, this feature would become less and less useful because of security considerations. However, this evolution may take a significantly long time. Because of the memory/software based SENET-DAX concept, it may be possible to take advantage of the increased capacity made accessible by TADI, within the constraints of the crypto environment, without unduly limiting the system.

1.3.3 <u>Routing</u>

1.3.3.1 Conventional Routing Methods

As indicated in Section 1.1, a conventional Space Division Multiplex (SDMX) system must commit (route and allocate) a fixed portion of its capacity to a potential call in order to establish whether such a call can be completed. That is, the established connection is used for the transmission of routing and signaling information between originating, intermediate, and terminating switches. For the time the calling party finishes signaling the originating switch until the receipt of an answer-back indicating that the called party has gone off-hook, the allocated channel carries only sporadic, low-information-content signaling messages. The remainder of the full-capacity channel is wasted, inaccessible to any other pending call (subject to the precedence scheme, of course).

An analogous situation occurs in Time Division Multiplex (TDMX) systems, where the allocation on behalf of a potential call is ordinarily a fixed increment of the transmission system's resources, consisting of every n-th slot in an n-channel TDMX. While from a technical standpoint, these slots need not be committed prior to the called party going off-hook, allocations are usually made when signaling is completed. Interswitch signaling in a TDMX system is frequently accomplished over a common channel; that is, all interswitch signaling is accomplished on behalf of all calls via one assigned TDMX channel. Establishment and maintenance of this Common Channel Interswitch Signaling (CCIS) channel reduces an n-channel TDMX to a capacity of n-1 channels available for subscriber traffic, even though instantaneous CCIS traffic rarely peaks near maximum channel capacity, and long term CCIS traffic is typically a small portion of the available capacity of a full TDMX channel. Thus in a typical

TDMX system, a certain amount of transmission capacity is wasted by allocation of fixed-size increments of transmission capacity to each call, regardless of the call's actual requirements (in fact, TDMX systems demand that each call be matched to the transmission rate of the channel, by bit-stuffing if necessary), and another portion of the system capacity is wasted because the unused portion of the CCIS channel capacity is inaccessible to subscriber traffic.

1.3.3.2 Advantages of Routing Under Dynamic Allocation

The SENET-DAX dynamic channel allocation capability eliminates the waste ordinarily incurred by the traditional SDMX and typical TDMX transmission systems. Unlike conventional systems, the SENET-DAX can allocate to each call precisely the transmission capacity that call requires. Transmission of CCIS messages as variable length Class II packets, of no greater length than required for the particular call function and transmitted only when necessary, results in optimizing both the CCIS response time and efficiency of use of transmission resources. CCIS packets are transmitted through the system as high-speed query/response type messages, yet transmission capacity is required only when a particular call function results in the queueing of a CCIS packet. In that event, the capacity of each link is utilized only for a single frame for each CCIS message. At all other times, that portion of the master frame used for Class II transmission is fully available for data packets.

A more subtle advantage of the SENET-DAX dynamic channel allocation concept over SDMX systems, and over most TDMX systems as presently implemented, is the capability in the SENET-DAX concept to separate the call connection process into two separate and independent phases: call reservation and call allocation. By relying on the very rapid transit of CCIS messages handled as Class II packets through the network, there is no necessity to allocate Class I system capacity until the called party has answered. Thus, from the time the calling party finishes signaling the originating switch until the answer-back from the called party reaches the originating switch, only Class II CCIS messages are required. Compared to immediate allocation of channel capacity in an SDMX or in most TDMX systems, this permits channel capacity to be assigned to other traffic less the CCIS messages for the given call. For example, assume that the average Class I allocation is 160 bits per master frame, the master frames have a 10 millisecond period, four CCIS messages flow in each

*It is shown in Section 10.2 that less than 1 percent of channel capacity is required for CCIS traffic.

direction per link during signaling, and the average subscriber delay in answering is 12 seconds. Then for CCIS messages averaging 160 bits in length and Class II data messages averaging 7420 bits in length, the SENET-DAX can transmit over 25 average length messages while waiting for the called subscriber to answer.

1.3.3.3 Effect on Class I Routing and Connection

The SENET-DAX efficiencies above those possible in conventional systems affect the routing of Class I traffic in the following way. For a conventional SDMX or TDMX system, transmission and handling of CCIS messages is often slow, usually being based on the equivalent of 2400 B/S transmission. In addition, route reservation and allocation of resources occur simultaneously. This often results in selection of a routing strategy based as much on meeting maximum connection time requirements as on optimal routing of the call through the network. In the extreme, for example, this can result in unnecessary all-trunks-busy rejection of calls for which the network as a whole has capacity, but which capacity cannot be found and connected within the time permitted. By contrast, the SENET-DAX concept routes the CCIS messages for call reservation along the route the call is expected to take, but only if that route is currently available and is optimum for Class II data packets between the two nodes involved. These CCIS packets can be alternately routed, if necessary, automatically and on a moment-by-moment optimum basis. Answer-back and call termination messages, and the SENET-DAX unique reallocation CCIS messages, will in general, be Class II routed completely independently of the Class I call route.

This advantage of the SENET-DAX dynamic channel allocation scheme can be utilized to permit quasi-associative routing of Class I calls in a dynamically optimized fashion as described for Class II traffic in Section 3.2. This is expected to result in call connection times equal to or better than those found in the best-performing conventional systems, combined with the advantages of a sophisticated routing scheme that optimizes utilization of transmission system resources.

1.3.4 Service Features

1.3.4.1 Range of Potential Service Features

The most significant impact of the SENET-DAX concept on service features is that dynamic channel allocation as structured in a stored-program distributed architecture permits an unprecedented degree of flexibility and efficiency in handling a variety of subscribers and trunks differing in many characteristics. These include urgency of connection/response time, security, precedence needs, digitizing algorithms, need for multi-party communications, error-tolerence, and others. It also provides flexibility for system necessities such as automatic alternate routing, intercept, preemption, journaling, retrieval and tracing.

1.3.4.2 Class I Service Features

Switches based on the SENET-DAX concept can feasibly offer to Class I subscribers all of the service features commonly available in a high-capacity conventional circuit switch. These features may include such capabilities as precedence and preemption, conferencing, call transfer and forwarding, compressed and abbreviated dialing, secure call mode conversion (see Section 1.3.4.4), automatic line/trunk grouping and hunting, direct access service, attendant recall/operator functions, and virtual connection of bulk data calls (such as sensor data and long-duration computer interconnections) and forward error corrected facsimile and low-speed video.

The impact of the SENET-DAX concept most visible in Class I features will be that subscriber equipments will not be constrained to a few types of loop appearances or a compatible family of appearances, but a wide range of equipments may receive service from the SENET-DAX. Less visible to an individual subscriber, but of obvious significance on a network basis, is the higher total throughput and improved network grade of service possible because of the dynamic channel allocation process. The subscriber will also see improvement in the speed of call connection and of data transfer through such a network.

1.3.4.3 Class II Service Features

Switches based on the SENET-DAX concept can also feasibly offer all of the service features commonly available in high-capacity conventional store-and-forward message switches and fast-response-time/high capacity "conventional" packet switches.

These features may include such capabilities as message accountability, message and communications line security, precedence and preemption, multiple addressing, universal data subscriber interfacing, mode/code/speed/format conversion, archival storage, intercept, retrieval and trace, bulk data transfer capability, automatic-repeat-request (ARQ) facsimile, including facsimile store-and-forward and facsimile multiple address, and rapid query/response, interactive, and data base update services.

The impact of the SENET-DAX concept most visible in Class II features will be that subscriber equipments of not only different characteristics, but of quite different demands for quality and speed of service, will receive such service simultaneously and non-interferingly. Less visible to an individual data subscriber, but of obvious significance on a network basis, is the higher throughput capacity and faster response/transit time possible for data messages traversing the network. Again, this is a direct consequence of the SENET-DAX dynamic channel allocation process. In the case of data transmission, it is calculable that total data transmission requirements will approximate 5 percent of the total voice transmission requirements. In essence this means that a network adequate in total capacity for projected voice transmission requirements during peak periods will normally offer data a much greater information transmission capacity than the peak data transmission requirements could justify when considered alone.

1.3.4.4 Digital Voice Conversion

Another interesting potential service feature brought to light during the course of this study is the conversion of voice calls between subscriber digital terminals that are incompatible in conversion algorithm, speed, synchronization technique, or any combination of these. The SENET concept does not directly make such conversions feasible. It simply establishes a communication system environment in which incompatible voice calls can exist simultaneously. Given this capability, the question arises whether the subscriber should be limited to communication only within a family of compatible instruments. In the case of data terminal equipments, the answer has for a long time been that mode/code/speed/format conversion was to be supplied in order to permit intercommunication among incompatible data subscribers. The question has not arisen significantly for voice networks. Terminal instruments tend to be compatible in an analog world, and only signaling compatibility has to be ensured. In a digitized voice environment, it has long been accepted that all subscriber instruments would have to be compatible.

Within the security constraints previously discussed (Section 1.3.2), which apply equally to data terminal equipment interconnection, and given development of practical algorithm-to-algorithm mapping techniques, the SENET concept could rather uniquely make effective use of such a conversion technique. This uniqueness is not so much in the provision of such a capability (this would be clearly desirable in any network tasked with interconnecting incompatible voice subscribers; for example, bit stuffing for speed compatibility between TDMX subscribers), but rather in using the technique to enhance transmission efficiency through the network. This does not appear to be true of either SDMX or of TDMX multiplexing.

Enhancement of SENET-DAX network transmission efficiency could be obtained whenever the voice subscriber terminal instruments operate at different transmission rates, and the call is either clear text or Crypto Stage I (link-by-link encryption). The improvement would be accomplished by always applying the conversion between the dissimilar rates as near as possible to the higher speed subscriber. For the interconnection of a 32 kilobits/second CVSD equipment and a 9600 bits/second APC equipment, for example, the conversion would be done in the SENET-DAX on which the CVSD subscriber was homed. Consequently, the call would be carried through the SENET-DAX network (in both directions) at 9600 bits/second, requiring only 30 percent of the transmission bandwidth necessary for a 32 kilobits/second call. Note that this approach would apply regardless of which subscriber initiated the call. The flexibility of the SENET concept also applies if the incompatibility between two digitized voice subscriber instruments is not one of speed, but one of technique or conversion algorithm.

The rapidity with which CCIS is accomplished can also lead to enhancement of system performance. For example, if conversion is required, and one of the SENET-DAX switches involved has already busied all of its conversion units, the call need not be denied. Instead, the system can arrange via CCIS messages to transfer the conversion process to a tandem SENET-DAX, or to the SENET-DAX at the opposite end of the connection, if it has a conversion unit available.

Finally, it should be noted (as pointed out in Section 1.3.2) that even when end-to-end encryption (Crypto Stage II) is common, a slight variation in COMSEC doctrine would still permit interconnection between incompatible subscribers by means of algorithm conversion. The variation (Stage I/II) would divide end-to-end encryption into two parts: from the originator to the converting SENET-DAX, and from the converting SENET-DAX to the recipient.

1.3.5 Network Management

The most significant impact of the SENET-DAX dynamic channel allocation concept on network management is that it permits extremely flexible and efficient handling of the interswitch communications necessary for the provision of alternate/adaptive routing, traffic load control, and network technical control functions for an integrated voice and data communications system. These are discussed briefly in this section. In addition, the possibilities of Time-Assigned-Speech-Interpolation (TASI) and Time-Assigned-Data-Interpolation (TADI) are considered in some detail.

1.3.5.1 Management of Adaptive Routing

SENET-DAX capability for extremely short delivery delays for CCIS packets handling virtual real time connections permits the improvement of transmission efficiency prevously described. A less obvious impact is that routing of CCIS messages can be accomplished in a quasi-associative manner, thereby enhancing SENET-DAX selection of the most effective alternate route. For example, to determine whether it is worthwhile preempting a link from A to C, system capability can be utilized by routing CCIS messages from A to B to C requesting confirmation of availability of the route from C on to the destination D. By contrast, conventional systems can usually meet connection time requirements by blindly preempting each successive link in sequential order, even if a non-preemptable link is eventually encountered. In this case, not only is the new call ultimately blocked, but other calls have been uselessly preempted.

Similarly, it is suggested that SENET-DAX routing of Class II Store-and-Squirt packets (Section 1.3.1.3) be adaptive, modeled on the ARPANET routing algorithm. This algorithm adjusts the routing from origin to destination by choosing the lowest observed delay route between each node and the next towards the destination. The SENET-DAX can utilize a similar algorithm, but need not base its operation purely on observation. During periods of relatively light traffic, the SENET-DAX could actively test transit times to other nodes with which it has not recently been in communication, using CCIS messages on a space-available basis. This would permit optimization of adaptive routing tables on a continuous basis, rather than occasionally handicapping the transit times of initial packets over an infrequently used route because of reliance on non-current information.

1.3.5.2 Management of Traffic Load Control

The speed and efficiency with which interswitch signaling can be accomplished by the SENET-DAX network should permit traffic load control to be done in a modulated fashion. Rather than each node either permitting all or no traffic entry, CCIS messages can be used to accept certain precedences of Class I traffic, while blocking or alternately routing others. Simultaneously, certain precedences of Class II traffic can be accepted, others delayed, and others routed alternately. These decisions can be made at each SENET-DAX, which knows not only the status of its own links directly to other SENET-DAX nodes, but also knows, when necessary, the status of remote links farther on in the network. For example, in the case of a node outage, rerouting could be accomplished at switching centers one remove back from those directly connected to the inoperable node. This should have the effect of spreading out and thereby minimizing the impact of such an outage.

1.3.5.3 Management of Systems and Technical Control

The flexibility of the SENET-DAX concept should permit very effective network technical and system control and management. In particular, the SENET concept is intended to be essentially transparent to many types of traffic, while affording them extremely high performance. This capability should permit economical technical and system control communications for relatively infrequent reports and commands, yet permit very fast dissemination and response times.

1.3.5.4 Time-Assigned-Speech-Interpolation

An evaluation has been conducted during the study of the practicality of making use of silent periods during transmission of digitized voice calls for transmission of other types of information. Silent periods occur during voice transmission when one of the parties is listening, and as pauses while one party is speaking (the former tending to be considerably longer). In either case, the silences represent periods, which for a typical SENET-DAX may be from a few to many master frames, during which the SENET-DAX has nothing to transmit in the allocation made for that voice call in the Class I region of its master frame. In current systems, techniques have been devised for detecting these silent periods and making use of them for other active calls. One of these techniques is called Time-Assigned-Speech-Interpolation (TASI). One particular part of the study was concerned with the practicality of such a technique in a SENET-DAX system.

Upon examination, it quickly became apparent that application of a TASI-like technique in the SENET-DAX network would be sensitive to the crypto environment. That is, detection of silent periods in near-real time could be accomplished (and in fact would be much easier with the digitized voice which the SENET-DAXs are specified to handle), only if voice transmissions are unencrypted or link-by-link encrypted (Crypto Stage I). But even in this case, the dynamic allocation capability of the SENET-DAX concept militates against use of these silences for speech interpolation. Because the makeup of each frame is based on a map of the Class I region, rather meticulous control of the alteration of these maps is considered necessary. Transmitting and receiving SENET-DAX must agree exactly on what is located where, and how much of the frame each allocation requires. Just as importantly, they must agree on the exact timing of constantly changing allocation maps. The resulting coordination complexity makes it impractical to interpolate speech into the silent periods of other speech.

In the case of end-to-end encrypted calls (Crypto Stage II), it would be impossible to detect silent periods in either direction and interpolation of any kind is not feasible. Unlike the situation of conversion between digitized voice techniques (Section 1.3.4), Crypto Stage I/II (end-to-end encryption to and from a single common point) would not be any more feasible than Crypto Stage II with respect to interpolation.

1.3.5.5 Time-Assigned-Data-Interpolation (TADI)

Although TASI techniques do not appear feasible in the SENET-DAX concept, interpolation of packet-switched data (TADI) into the silent periods of clear voice or Stage I encrypted voice appears a possibility. Consider the following hypothetical technique. Immediately following the Start-of-Frame marker, each master frame could include a Frame Silences Map (FSM). This map would contain one bit position for each of the maximum number of Class I calls a master frame could contain. For each Class I connection in a particular frame having silence, the FSM would contain a zero-bit; for each connection in the frame having voice at the moment, the FSM would contain a one-bit. Any unassigned allocations (including those beyond the normal Class I/Class II boundary marker) would always have zero-bits at their positions in the FSM; Crypto Stage II calls would always have one-bits at their positions in the FSM.

The SENET-DAX would treat all silent allocations in each frame as a dispersed string of sequential bit spaces, followed contiguously by the formal Class II region. Using normal procedures, the SENET-DAX would start with the first packet of the oldest message of the highest precedence, and would allocate the bits of this packet sequentially into as many of the silent allocations in the Class I region of that frame as are necessary. When a complete packet has been mapped into the region of silences, the next queued packet in order will be similarly and contiguously mapped into the region of silences. If the bits available in the totality of silent allocations in a frame were not enough to contain all of the packets queued and ready for transmission during that frame, the Class II region would be used as a contiguous extension of the silent allocation.

On receipt of a master frame in which data packets were being interpolated into silent periods of Class I digitized voice allocations, the receiving SENET-DAX would utilize the received FSM to determine which allocations in that particular frame should be handled as normal Class I digitized voice calls, and which (silent) allocations would be string processed as sequential dispersed bits, with the addition of the literally contiguous Class II region bits if necessary.

Analysis based on DOD projections of voice and data traffic during the 1980's, including typical digitized voice and data rates, indicates that a typical SENET-DAX master frame, during a peak second in this time period, would comprise approximately 82.8 percent Class I digitized voice, and 17.2 percent Class I non-voice (facsimile and slow-scan video) and Class II data traffic. These same traffic analyses indicate that in a hypothetical network projected for this time period, a single T1 carrier may not suffice for network links in some cases. However, for illustrative purposes. consider utilization of a single T1 carrier link using TADI of Class II packets into Class I digitized voice silences. During each peak second, a SENET-DAX utilizing a 1.544 megabits/second T1 carrier to capacity without consideration of TADI (Time-Assigned-Data-Interpolation) would transmit in each 10 millisecond master frame a 7-bit Start-of-Frame marker, approximately 12,773 bits of Class I digitized voice (82.8 percent of net), a 7-bit Class I/Class II marker, and approximately 2,653 bits of non-voice Class I and Class II data traffic. In other areas of this study, a weighted average voice allocation was calculated to be approximately 155.4 bits per 10 millisecond master frame; thus voice allocations during the peak second equate to approximately 82 concurrent voice calls. In order to utilize the TADI technique

hypothesized, the Frame Silences Map must provide for the maximum number of voice calls the master frame need handle. Nominally, this would be 643 (assuming an entire master frame of 2400 bits/second vocoder calls), as opposed to the 99 which would be required if the master frame were fully occupied by average length digitized voice calls. Arbitrarily (and approximately) splitting the difference, it is likely that a 384-bit Frame Silences Map will suffice. Assuming that this 384-bit FSM is transmitted immediately following the Start-of-Frame marker, and that the master frame is still handling the same 82 average length digitized voice calls, there would be left 2306 bits for non-voice traffic in the peak second. Of these bits, approximately 719 would be required for the non-voice Class I traffic (slow-scan video and FEC facsimile), leaving some 1587 bits to be used for Class II data traffic.

Studies made in conjunction with the TASI technique indicate that an average two-way four-wire connection would be active (i.e., non-silent) about 35 percent to 40 percent of the time during any given frame, allowing for listening periods and for silences during speech. Therefore, for the 82 voice call allocations, some 49 to 53 allocations, totaling 7,615 to 8,236 bits, would be available for Time-Assigned-Data-Interpolation during a peak second in each master frame, exclusive of the nominal Class II region. Since the original 2,653 bits of non-voice Class I and Class II data traffic include 719 bits of slow-scan video and FEC facsimile, there remain some 1,934 bits of Class II data traffic to transmit in the master frame during the peak period. From 3.9 to 4.3 times this many could be inserted into the silent portions of the original 82 digitized voice connections, if the hypothetical TADI technique were used. To this capacity could be added the nominal Class II region (1587 bits) for even more capacity, if this were necessary. The resulting total capacity for Class II data, concurrent with the original 82 digitized voice calls and with the calculated amounts of slow-scan video and FEC facsimile, would be from 9,549 to 10,170 bits.

It appears, however, that for a SENET-DAX network which handles digitized voice traffic that is preponderantly clear or link-by-link encrypted (Crypto Stage I), a better use of the TADI technique would be to expand the digitized voice capacity of the network. Allowing the 7-bit Start-of-Frame marker and a 384-bit Frame Silences Map, as above, and allowing the Class I traffic to fill the rest of the peak second master frame, a typical master frame would include 15,049 Class I bits, or some 11.5 percent more than before. Of these bits, approximately 802 would be non-voice Class I

(slow-scan video and FEC facsimile) and not subject to data interpolation. The 14,247 Class I digitized voice bits would correspond to approximately 92 average length voice allocations, 10 more than before. These 92 voice allocations would in turn include, in an average master frame, from 55 to 59 silent allocations, totaling 8,547 to 9,168 bits available for TADI of Class II data packets. It appears therefore that utilization of this TADI technique would, under peak second conditions, permit an increase of 11 percent to 12 percent in the amount of Class I traffic a SENET-DAX network could handle, while concurrently the amount of Class II data traffic that could be handled would be from 4.4 to 4.8 times as great as could be handled without Time-Assigned-Data-Interpolation, barring excessive amounts of Crypto Stage II digitized voice traffic.

The TADI approach is promising but, because of a significant software impact, has not yet been explored in depth.

1.3.6 Subscriber Terminal Requirements

Although study of the SENET-DAX concept at this point is still an investigation of techniques, it is nontheless useful to begin consideration of the technical impact of the dynamic channel allocation concept on subscriber terminal requirements. Investigations to date indicate that there appears to be no direct impact of the SENET-DAX concept on the requirements for subscriber terminal equipments homed on a switching center. The only immediately evident indirect impact appears salutary; the SENET-DAX concept seems well suited to integrating efficiently and flexibly the transmission of communications between members of widely disparate subscriber communities.

1.3.6.1 SENET Concept for Loops

As far as has been determined during the course of this study, it appears improbable that the SENET concept would ever be used on the loop (subscriber) side of a SENET-DAX. It is not anticipated that subscriber terminals would ever exhibit varying types of data communication aspects in dynamically varying proportions. Rather, it seems more probable that any individual subscriber terminal would present an invariant aspect to the system at any one time.

As a corollary to the foregoing, there does not appear to be any particular usefulness in development of a subscriber terminal capable of generating its allocated portion of a SENET-DAX master frame at the rate required for trunk transmission, and with the coordination necessary for smooth integration into a master frame with other portions of that master frame generated internally to the SENET-DAX. To begin with, this would require loops capable of the same transmission rates as the trunks. This currently would be quite uneconomical. Of considerably more import, the coordination of dynamic channel allocation would become very unwieldy, and probably quite inefficient. Finally, the whole point of the SENET concept of software-controlled multiplexing to and from a shared memory would be vitiated by such direct connections.

1.3.6.2 Line Termination Units

Current communications switching systems ordinarily require line termination units to facilitate the interfacing of subscribers on individual loops with the multiplexed transmission trunks. This is especially necessary when an analog subscriber is to be multiplexed onto a digital trunk. Similarly, since conventional store-and-forward and packet switching data communication systems operate into and out of memory under software control, line termination units are required to interface widely dissimilar types of subscribers with a shared memory. The SENET-DAX has requirements similar, and of no greater or lesser complexity, to these conventional systems. A SENET-DAX on an individual basis might require a greater variety of line termination units than any single-subscriber type of communications switch now available, but this is a simple consequence of the flexibility of the SENET-DAX concept.

SECTION 2

MASTER FRAME STRUCTURE

SECTION 2

MASTER FRAME STRUCTURE

2.1 FRAMING AND ENVELOPING

2.1.1 Problem

The transmission concept we are studying is an integrated switched telecommunications network which is to handle simultaneously, circuit switched, packet switched, and some message switched (Store & Forward) traffic. The concept is expected to permit a high level of transmission efficiency not obtainable by separate systems. The idea is based on the dynamic allocation of variable sized increments of the transmission capacity of a master frame period as a function of traffic demand by class. This section will attempt to define the structure of that master frame and the framing/envelope characteristics of the various classes of traffic served by the system. This structure, along with factors affecting the selection of frame period, will be analyzed and redefined, if necessary, to minimize cross-office and total cross-network time delay and impairment of system performance.

2.1.2 Objectives

We are interested in achieving the following objectives as a result of our solution of the framing/envelope problem.

1. To support subscribers using equipment characterized by digital data rates in the 8000N family, the 2 x 75^N family, commercial common carriers, and transmission systems now under development.

2. To handle digitized voice rates representative of various evolutionary stages of development simultaneously with the data rates above without mutual interferences or prohibitive inefficiency due to catering to the least efficient technique.

3. To improve the grade of service, increase the capacity, or both, as more efficient voice digitization, facsimile transmission, and video transmission digitization techniques are

introduced. These improvements should be independent of each other.

4. To minimize the cross-network delays and timings involved in establishing calls, coordinating connections, crypto-synchronization, propagation of trunk signalling (CCIS) messages, error control and channel coordination in noisy or delay environments, and end-to-end control and accountibility effects.

5. To minimize cross-office network delays and timings for implementation concepts for DAX's and integrated tandem switches.

2.1.3 Analysis and Results

2.1.3.1 Introduction to the Master Frame Format

Our approach will be to divide wideband digital trunk capacities into master frames in the time domain (see Figure 2-1). The period of the master frame, once selected, would be constant (e.g. 10 msec or 20 msec) throughout a system. However, within a frame, traffic of varying bandwidths can be accomodated simultaneously by allocating time shares as needed, until the total capacity of the wideband trunk has been allocated. Within a master frame, portions will be utilized as a Start of Frame (SOF) Marker, Class I Region, and Class II Region, where CCIS messages represent a subset of the Class II allocation. Our choice of formats for the Class I and Class II Regions, and the CCIS Subregion, is shown in Figure 2-2.

The Master Frame Format must be compatible with terminal equipment having rates characteristics of the Defense Communications System (DCS), commercial common carriers, allied communications systems (e.g. NATO), and transmission systems under development. There is also a requirement to support subscribers utilizing equipments characterized by digital data rates in the 8000N family, the 2 x 75^N family, and at rates not

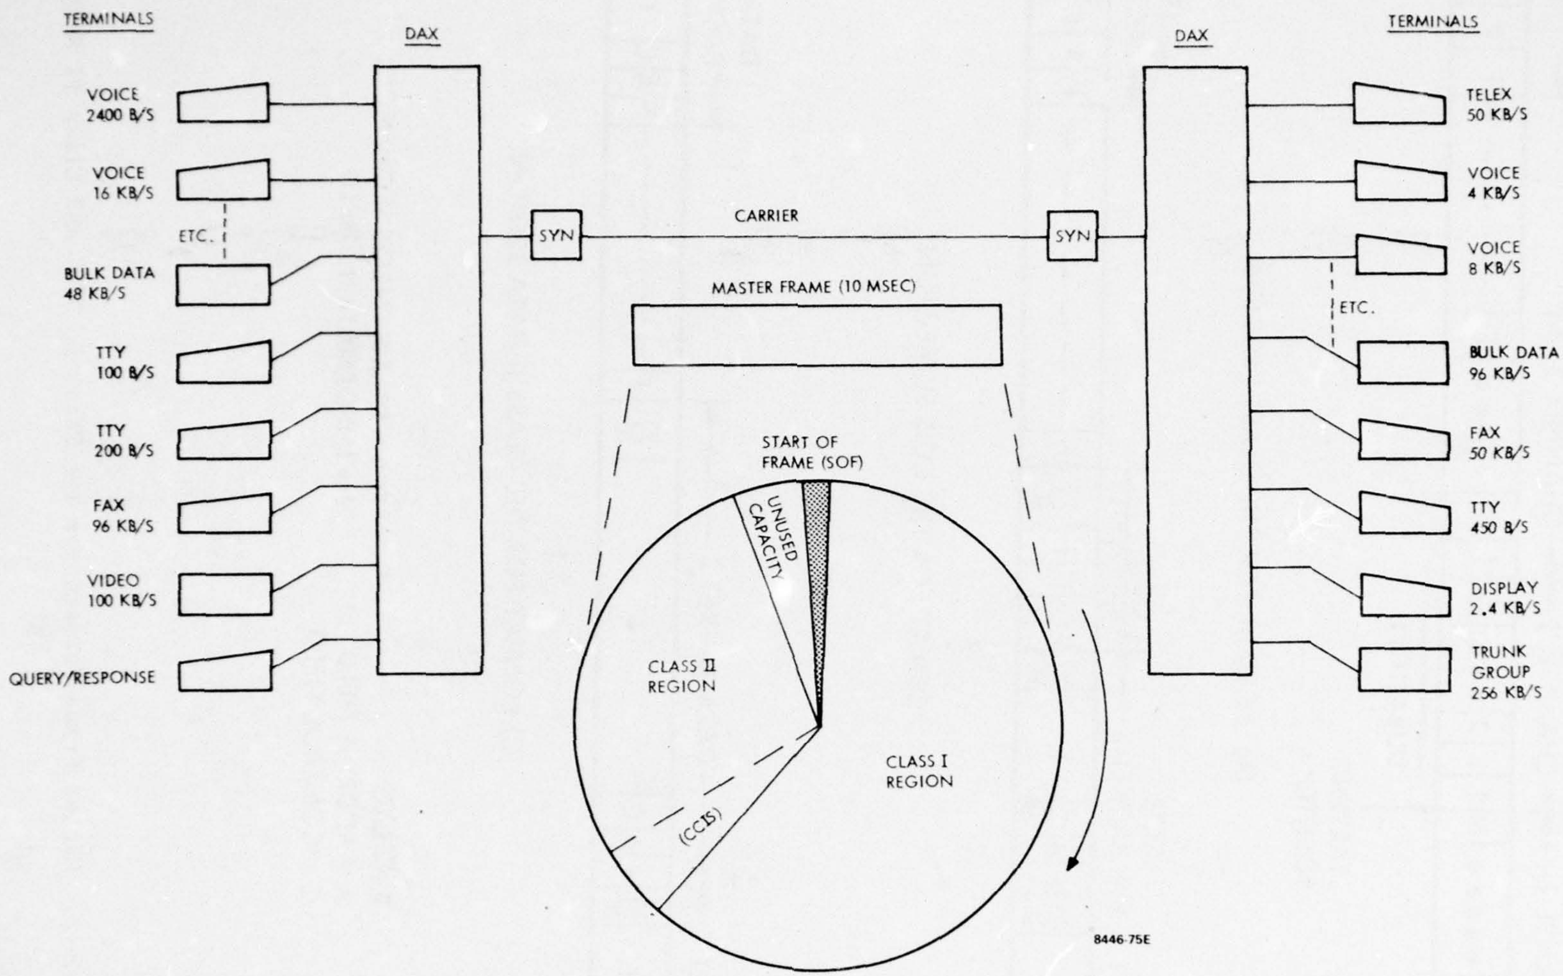

Figure 2-1. Allocation of the Master Frame on a DAX Link

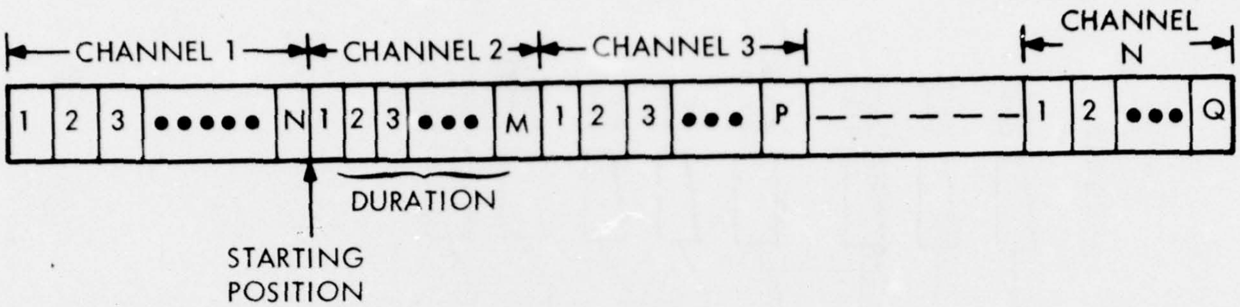

(A) FORMAT FOR THE CLASS I REGION

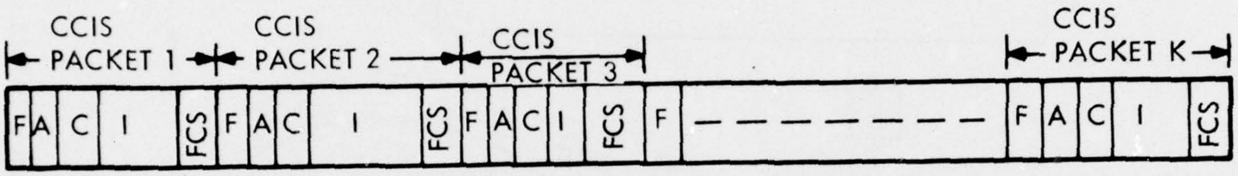

(B) FORMAT FOR THE CCIS SUBREGION

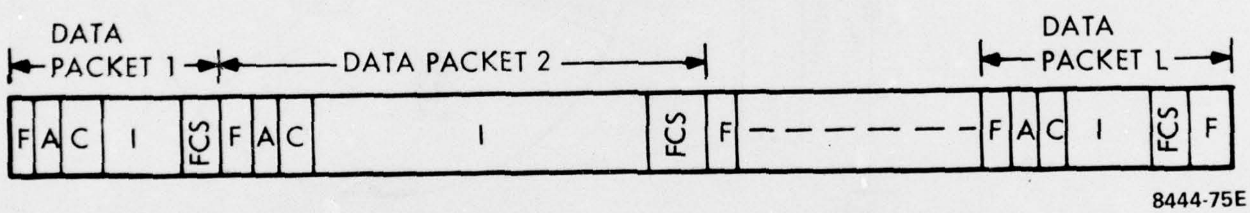

(C) FORMAT FOR THE CLASS II DATA REGION

F = FLAG
A = ADDRESS FIELD
C = CONTROL FIELD

FCS = FRAME CHECK SEQUENCE
I = INFORMATION FIELD

Figure 2-2. Master Frame Formats for the Class I, CCIS and Class II Regions

found in either family. Because Class I calls must be processed in real-time, while Class I packets will be speed-buffered, the period for the master frame will be chosen primarily as satisfactory to Class I traffic. Class II traffic will be handled within this constraint.

2.1.3.2 Start of Frame Marker

The Start of Frame (SOF) Marker is used for synchronizing at the beginning of the Master Frame. It consists of the first N bits of a frame and is expected to take on one of two possible forms. The first is a relatively long N-bit coded sequence or a series of short N_2-bit coded sequences used to establish initial frame synchronization or to reestablish synchronization after a loss of sync. This is a powerful technique, verified to permit achievement of synchronous operation under the worst conditions anticipated for the application of this technique. The number of bits N_1 required in this case should not exceed 100; N_2 is on the order of 15 bits. The second form of SOF would be a short control sequence on the order of 1 to 15 bits, which would be used to monitor and maintain sync during the time when both transmission directions of a link are in-sync.

The adaptability of this scheme is attractive from the point of view that it minimizes transmission overhead and makes use of the considerable software capability of the DAX. Use of the single short control sequence would seem to provide the most efficient frame overhead utilization; however,this results in a less rapid sync acquisition - too high a price to pay for the increase in efficiency. Section 4.2 - Frame Synchronization deals with the Start of Frame Marker in some detail.

2.1.3.3 Class I Region

The Class I Region will be utilized for transmission of Class I circuit switched traffic, to include digitized voice, facsimile, and low speed digital video. Transmission rates utilized for voice will be those of the 8000N family (4 kbps, 8 kbps, 16 kbps, 32 kbps) and of

the 2 x 75^N (2400 bps, 4800 bps, 9600 bps) as well as others - (50 Kbps). The range of speed for video communication will be 100 Kbps to 200 Kbps and for facsimile will be 50 Kbps to 64 Kbps.

The format of Class I traffic is similar to that of a time-division demand assignment approach. Because Class I is defined to comprise real-time traffic, the time slice to which each call is assigned is reserved during succeeding frame periods until the termination of the connection, or until termination of a call placed earlier permits reallocation of an on-going call. Time slices are allocated and controlled through signalling and supervision messages in the CCIS Sub-region. The main differences between this approach and that of standard synchronized TDM are that allocations of time-shares are variable to accomodate Class I traffic at different bit rates and the location of these allocations in the master frame are not fixed, nor is the total information bandwidth limited (except by the capacity of the trunk).

The time structure of the Class I Region is shown in Figure 2-2a. With a 10 msec master frame as derived from a 1.544 Mbps T1 Common Carrier, a total of 15,440 bits is provided in each of 100 master frames per second. Digitized voice at 8 Kbps would then require an 80 bit allotment per frame. Taking this as the first channel in Figure 2-2a, bits 1 through M = 80 would be allocated for the call. A second channel at 2400 b/s would be allocated the next NL = 24 bits, etc. Each channel is uniquely indentified in the frame by its starting position and its duration in bits.

The advantage to this approach is that the system can handle simultaneously many different digitized voice rates representative of various evolutionary stages of development without mutual interferences or inefficiency due to catering to the least efficient techniques. Also, as more efficient techniques of voice digitization, facsimile transmission, and video digitization are introduced, these improvements may be introduced independent of one another and without changing the way they are handled in the network.

2.1.3.4 CCIS Messages

Following Class I traffic and occupying the first portion of the Class II Region will be the Common Channel Interswitch Signalling (CCIS) messages. CCIS messages will be sent in packet form according to the standard specified by Advanced Data Communication Control Procedures (ADCCP)*. This structure consists of a flag which is the 8 bit pattern 01111110, followed by an 8 bit address and 8 bit control field. A variable information field follows the control field and a 16 bit redundancy field for error control precedes the closing flag. The flag which closes a packet may also be used as the opening flag for the next sequence.

CCIS messages will contain information necessary to establish and route Class I calls through the network. CCIS and Class II data messages are contiguous and indistinguishable from each other except by an identification bit located in the address field. The choice of packetizing data using ADCCP format was made because it is very efficient for transmitting Class II data. The same procedure is used for CCIS, even though the overhead for the format is quite high. The reason is that the convenience of using the same format throughout the Class II Region far outweighs the overhead inefficiency involved with using ADCCP procedures for CCIS.

The CCIS region will be used to map the Class I allocation changes of subsequent master frames at the earliest. This is because processing demand is considerably eased when connection time is extended by one or more frames rather than required in the same frame period as the CCIS information transmitted. Thus the CCIS Region is positioned in the frame after the Class I data since it refers to changes made in frames subsequent to it. Another advantage of doing this is that CCIS and Class II data are both arranged in ADCCP packets and from a processing point of view they can be treated as one class. Thus the processing for voice and packets can be done separately and the boundary identifed

---

*For further information on ADCCP, see Section 2.4-CCIS fields and Formats.

by the flag of the first CCIS packet following the Class I traffic.

It is not expected that many CCIS messages will occur during the duration of individual Class I calls. CCIS messages are sent to originate calls, terminate calls, and reassign channel allocations to fill in gaps due to termination of other calls. Since many thousands of frames, e.g., 30,000 frames assuming an average holding time of 5 minutes and a frame period of 10 msec, will occur within the duration of an average call, most of the time no CCIS information will be sent. This factor places an upper limit on the CCIS Region which will be necessary to describe Class II traffic.

2.1.3.5 Class II Region

Following the CCIS message will be Class II data in the form of ADCCP packets,each self-identifying and self-routing and generally of variable size*. This provides support to data terminals which operate over a wide range of speeds. Incorporated in the packet format will be information necessary for routing, security, identity, and precedence. By keeping this information short with respect to the packet length, efficient utilization of the transmission efficiency of the system is possible.

Class II data traffic will include two categories; data characterized by short discrete messages requiring near real-time delivery (including interactive communications, query/response communications and data base updates) and bulk data characterized by long messages without the requirement for immediate delivery or on a FIFO (First-In, First-Out) basis by precedence level. Long bulk data messages with

---

*The packet size is variable within the limits of the ADCCP Frame Check Sequence (FCS) to provide adequate error protection for the particular environment encountered. For example, for Line-of-Sight Communications, the maximum packet size might be on the order of 1000 bits while for Tropo it might be only 400 bits.

typically low precedence and security classification, may be temporarily inhibited (probably at the terminal but possibly at other points in the system) when the demand on traffic is high and reinitiated when capacity becomes available. Justification for conversion of bulk data rate packets for Class II transmission is taken up in **Section 2.2- Packet Switched Bulk Data Transfers.**

2.1.3.6 Packet-Switched Class II Traffic

In this section we will consider a specific case which will illustrate the advantages of using ADCCP packets to transmit Class II data traffic. Figure 2-3 shows the structure of the two approaches we investigated to transmit Class II messages; ADCCP packets[3] and USASCII message blocks[38]. In ADCCP, signals for data link control are based on bits and unique bit patterns while for USASCII control signals are based on 7 bit ASCII code characters (with parity bit added).*

Both message types have overhead bits associated with them. For the ADCCP packet, overhead consists of the 40 bits which make up the Flag (F), Address (A), Control (C), and Frame Check Sequence (FCS) fields. This remains constant regardless of the length of the Information Field (up to the maximum length of the message). Overhead for the USASCII message includes 8 bits each for the Synchronization (SYN). Start of Message (DLE), Address (A), Control (C), End of Message (ETB), and Block Check Sequences (BCC) characters plus one parity bit for every ASCII character in the Information Field.** Thus the ADCCP packet is more efficient with respect to overhead than its USASCII counterpart. For the minimum message transmitted, the ratio of information bits to total bits transmitted is 40% for ADCCP while for USASCII it is approximately 30%. Information bits are defined here to be those bits not

---

*It is worthwhile to point out that the data communications industry has recommended that bit oriented ADCCP be adopted as the data link control standard for communications within computer networks.

**Each Address, Control, and Information Field character is composed of seven-bit ASCII characters with a parity bit added.

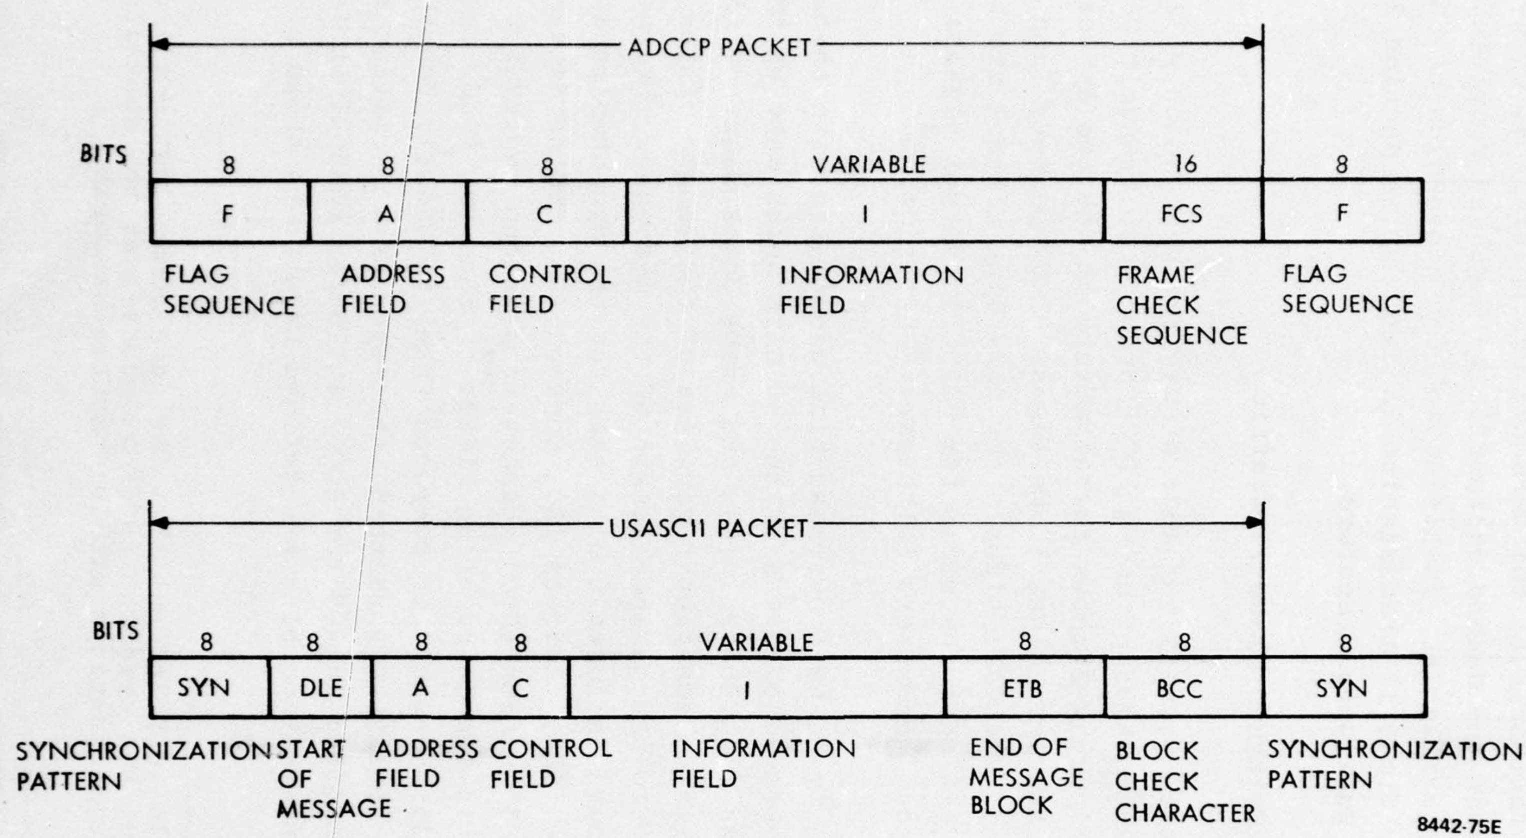

Figure 2-3. Format Structure of ADCCP Packets and USASCII Blocks

associated with either message delimiters or error correction. In Figure 2-3 this would include all the bits of the Information Field (I) plus those bits in the Address (A) and Control (C) Fields which concern message handling. For ADCCP format, the full 8 bits of the Address and Control Fields would be classified as information while for USASCII format only 7 bits would qualify (the 8th bit in each field is a parity bit).

Transmission efficiency for each message type as the Information field is varied is given in Table 2-1. For equivalent Autodin message transmission, i.e., when the text portion is composed of 80 - 8 bit information characters, ADCCP exhibits an overhead percentage (3.6%) less than half that of USASCII (9.1%).

TABLE 2-1. TRANSMISSION EFFICIENCY AS A FUNCTION OF INFORMATION TRANSMITTED

Information Bits	Transmission Efficiency	
	USASCII	ADCCP
0 bits	29.2%	40.0%
200	85.3	89.3
400	92.1	94.3
600	94.1	96.2
800	95.9	97.1
1000	96.7	97.7
1200	97.2	98.0

Besides being more efficient, ADCCP has more inherent descriptive power than USASCII. In ADCCP the full Address and Control fields are available for use while in USASCII, only the first seven bits are available (the 8th bit is for parity). Thus twice as many addresses (and control) can be specified using ADCCP as can be specified using USASCII.

A third recommendation for the ADCCP format concerns the error control capability. Checking the ASCII code of the USASCII variable length message for errors is done using horizontal and vertical parity bits, while for ADCCP polynomial codes are used. Polynomial checking is generally a more powerful technique than the ASCII horizontal and vertical checking, although the cost of encoding and decoding equipment is higher.* For example, using the 16 bit Frame Check Sequence for protection, and transmitting fixed length blocks of 800 bits over a telephone line with a random bit error rate probability of 10^{-3}, the probability of obtaining an undetected error on the block is of the order of 10^{-8} (see Reference 38). The corresponding undetected error probability for horizontal or vertical parity checking is of the order of 10^{-7} (see Reference 68). Thus ADCCP exhibits a higher degree of protection than its USASCII counterpart.

2.1.3.7 Packet-Switched Bulk Data Traffic

Although the practicality of using packets to transmit bulk data traffic will be fully covered in Section 2.2, we will briefly touch on the justification of this approach. Bulk data transmitted in message blocks stands a relatively high probability of being interrupted frequently through preemption, especially for long data records. Using ADCCP packets, we acquire a capability for long range transmission of bulk data packet by packet, with each packet being transmitted to its destination independent of one another, often along different routes,

---

*It should be noted that with LSI technology advances, the cost associated with circuitry to implement the more complex codes is becoming much more feasible.[38]

and in small blocks which are rarely, if ever, preempted. Even if they are preempted, only one packet is affected, not the whole message and this packet is only slightly delayed. At the receiving end, the packet can be reassembled off-line in storage according to sequence number to make up the original bulk message. This noninterruption of service leads to a practical method of bulk data transfer in a military preemptive environment and is the reason we chose this particular method.

2.1.3.8 Dynamic Interaction Between Classes

The Class I and Class II Regions will be dynamically assigned during successive master frames depending on the traffic demand and precedence of a call. Class I traffic is first allocated capacity, upon demand, for all calls up to the capacity of the Master Frame. The Class I Region will normally be a large portion of the Master Frame. Calls requiring capacity beyond full capacity allocations are blocked and do not receive service unless they are of preemptory precedence. After CCIS messages are transmitted, Class II interactive and bulk data packets will be sent in FIFO order by message precedence. Any portion of the frame not allocated to Class I traffic due to low Class I traffic intensity, will be used to transmit Class II data packets. This dynamic assignment of the Class I and Class II Regions will permit on a continuous basis the grouping of these traffic classes to minimize the inefficiency of having unused space within a master frame. It has been established by several investigations that a system which shares its transmission capacity in this manner among a mixture of traffic classes requires less overall capacity than separate systems for the equivalent performance.[19] The result is a system that trades off higher net transmission capacity with greater organizational complexity and memory, and it is believed, lower net transmission cost.

2.1.3.9 Precedence/Preemption

Both Class I and Class II traffic will feature various precedence levels. One problem which now arises is how to equate precedence of different types of traffic in order to determine the rules governing the preemption of calls. The present approach is to equate the precedence of

voice and data traffic on a one-to-one basis, i.e., a data call (Class II) of higher precedence can preempt a voice call (Class I) of lower precedence, if necessary, and vice versa. However, to ensure timely and efficient handling of Class I calls, the CCIS messages associated with these calls could probably be given a uniform, arbitrary precedence ensuring their propagation. One possibility is to assign FLASH precedence to Class I CCIS messages.

In the case where both voice and data calls have equal precedence and either call is preemptable, the data call will be preempted first. This is because preemption of a voice call implies a more catastrophic effect, since it involves breaking down the call and dropping it, whereas, a data call may be delayed and queued for transmission at a later time.

2.1.3.10 Slot Size

In Section 10.2 we will examine the efficacy of standardizing on a maximum transmission increment, or slot size, by which the bit position and the bit duration within a master frame for Class I calls could be specified. This approach eas assessed with respect to the efficiency of frame utilization in terms of the overhead bits needed to describe the positioning and allocation of a Class I call and the bit stufffing necessary to handle the wide variety of Class I data rates served by the system. Slot sizes of 1 and 4 bits for a frame period of 10 ms and 1, 4, and 8 bits for a 20 msec frame period appeear favorable from an implementation as well as a transmission efficiency point of view.

2.1.3.11 Connection Time and Performance

In Section 10.3, other factors are evaluated for their impact on cross-network timing and performance. These include accountability, dynamic alternate routing, crypto operation, multi-homing, buffering, and slot size as a function of efficiency. While connection times for traffic should be minimized (to a few seconds at most), these should be traded off against processing load (Section 3) and traffic handling requirements (Section 11). For example, it appears desirable to limit the total number of CCIS messages in a single master frame so as to bound the network against random surges of traffic, even at the cost of delay in handling Class I traffic.

2.1.3.12 Error Control

Error control and correction (see Sections 4.6, 9, and 10.2) depend greatly on the transmission path. A hybrid Forward Error Control/Automatic Repeat Request (FEC/ARQ) protect scheme adequate for the worst links (BER's of 10^{-3} or greater) turns out to be extremely inefficient for the best links. Automatic Repeat Request (ARQ), while powerful and effective for moderate error and short delay, falls apart under high noise and/or long path conditions, due to the frequency of retransmits. These and other procedures, such as adaptable approaches, require further evaluation in terms of capability and implementation complexity as function of postulated realistic noise and network environments.

2.2 PACKET-SWITCHING BULK DATA TRANSFERS

2.2.1 Problem

The SENET-DAX concept envisages handling in an integrated fashion the transmission of digitized voice, low-speed video, and Forward Error Correcting (FEC) facsimile via variable, dynamically-allocated virtual connections within a Class I region of the Master Frame. Common Channel Interswitch Signaling (CCIS) messages arranging for the establishment, reallocation, and termination of these virtual channels, and digital data transmissions involving interactive, query/response, data base update, and narrative record traffic are to be transmitted in those portions of Master Frames not required for Class I virtual connections, on a First-In/First-Out (FIFO) basis by precedence; these transmissions are classified as Class II in the SENET-DAX concept. Details of the various formats and procedures will be found in other sections of this study report.

In the SENET-DAX concept, very long data transfers of bulk data are categorized as Class III data transmissions. Appendix A to the SENET-DAX Study Statement of Work lists the following characteristics for Bulk Data:

a. Cross network connection time of 30 to 60 seconds;

b. Variable input data rates from 100 kb/s to 1 mb/s;

c. Error control is required;

d. Perogative of system to terminate call temporarily;

Other DCA-provided information (1) has been to the effect that bulk data transfers might be required by approximately 20% of the I/O terminals foreseen for the near future; that two different types of bulk data transfer can be foreseen, and can be conveniently labeled Bulk 1 and Bulk 2; that Bulk 1 transfers are expected to have nominal message lengths of 10,000 to 1 million bits, occur up to 16 times per busy

hour (8 transmissions and 8 receptions, both typically at 4800 b/s) at 19% of the total I/O terminals foreseen; and that Bulk 2 transfers are expected to have nominal message lengths of 1 to 100 million bits, occur approximately once in ten busy hours (primarily as transmissions, with negligible receptions, both typically at 9600 b/s) at 1% of the total I/O terminals foreseen. In this earlier DCA information, Bulk 1 was identified as being typically the transmission of entire files, programs, or processing results; Bulk 2 was identified as being typically the transfer of extremely lengthy information such as an entire data base or sensor data.

In order to analyse the practicality of packet-switching bulk data, some estimation of the magnitude to be transferred, the rate of transfer required, and the degree of tolerance of connection time and/or temporary interruptions must be formed, and combined with an estimate of the frequency of occurence of these transfers. In addition, procedural requirements (e.g. error control, channel coordination, accountability, store-and-forward requirements, etc) must be assessed. In addition, the demands to be placed on the anticipated capacity of the SENET-DAX utilizing realistic trunk groups, by other classes of traffic such as the non-deferrable Class I virtual connections must be assessed.

Table 2-2 summarizes the information drawn from both Appendix A to SENET-DAX Study Statement of Work, ard from the earlier DCA information [21] as applied to a system hypothetically based on utilization of T1 carriers having trunk group capacities of 1.544 millon bit per second information transfer rates. Table 2.2 highlights the virtual connection time for each Bulk data class at the transmission rates of interest, as well as the percentage of the nominal SENET-DAX capacity (based on the T1 carrier capacity) consumed for that virtual connection time when the Bulk data transfers are handled as Class I (virtual connection) calls, and when the transfers are handled as Class II (packet-switched) calls.

TABLE 2-2. BULK DATA TRANSFER DEMANDS ON T1 CARRIER-BASED SENET-DAX

TRANSFER TYPE	TOTAL BITS PER TRANSFER	TRANSFER (1) RATE (TYP)	VIRTUAL (2) CONNECTION TIME	T1 CARRIER UTILIZATION/FRAME CLASS I XFR	CLASS II XFR
BULK 1	$10^4 - 10^6$	4.8 KB/S	2.08 - 208 SEC.	0.31 %	0.32 %
BULK 1	$10^4 - 10^6$	100 KB/S	0.10 - 10 SEC.	6.48 %	6.74 %
BULK 1	$10^4 - 10^6$	1000 KB/S	0.01 - 1 SEC.	64.77 %	67.36 %
BULK 2	$10^6 - 10^8$	9.6 KB/S	1.74 - 173.6 MIN.	0.62 %	0.65 %
BULK 2	$10^6 - 10^8$	100 KB/S	0.17 - 16.7 MIN.	6.48 %	6.74 %
BULK 2	$10^6 - 10^8$	1000 KB/S	1.00 - 100 SEC.	64.77 %	67.36 %

NOTES: (1) SENET-DAX SYSTEM PEROGATIVE IS TO THROTTLE MESSAGE INPUT

(2) PLUS NETWORK CONNECTION TIME OF 30 TO 60 SECONDS/TRANSFER

2.2.2 Objective

The objective of this task is to determine the practicality of utilizing packet-switched envelopes for bulk data transfers. This requires consideration of the following points:

1. How may the permissible 30 to 60 seconds of cross network connection time be best utilized?
2. Basically, can the SENET-DAX transfer bulk data at rates up to 1 million bits per second?
3. Assuming a bulk data transfer rate range of 4.8 to 1000 KB/S, and a range of nominal bulk data message lengths of 10,000 to 100 million bits, is there a single optimum answer to this task?
4. What form of error control should be utilized for bulk data transfers?
5. How may the SENET-DAX system perogative of throttling (temporarily shutting off the message inputs) be utilized?
6. If bulk data is packet-switched, what relation should exist between the handling of Class III data vis-a-vis Classes I and II, and the CCIS messages?
7. What degree of accountability should the various SENET-DAX's processing Class III data afford that data?

2.2.3 Analysis and Results

Examination of the problem parameters stated in Section 2.2.1 has led to the initial conclusion that there is an answer to Question 3 of the Objectives stated in Section 2. Assuming that the problem parameters stated are adequately precise for this study, it appears that more than one technique for transmission of bulk data should be provided by the SENET-DAX, or alternatively, that the SENET-DAX transmission of bulk data should be constrained realistically. Consequently, the analysis has been directed toward two goals: determination of a range of capabilities adequate to handle the entire stated range of bulk data transfer requirements; and definition of a subset of the stated requirements which the SENET-DAX can handle optimally. This initial conclusion has formed the basis on which the results reported in this section were founded, and accounts for the quality of answers found in response to the questions stated as the objectives of this part of the study.

2.1.3.1 Utilization of Cross Network Connection Time

Establishment of a Class I virtual connection over which to transmit bulk data should take no longer than the nominal 5 seconds specified in Appendix A to the SENET-DAX SOW. To establish a Class II packet-switched path over which to transmit packets of bulk data should require no longer than establishment of such a path for interactive packet data transmission, specified in Appendix A as a nominal 30 seconds. Since Appendix A also permits the SENET-DAX the perogative of throttling input data, it appears that the SENET-DAX through which a bulk data transfer is being originated should utilize its relatively relaxed call connection time to establish the availability of the subscriber station which is to terminate the bulk data transfer prior to accepting any of the bulk data transfer. This is considered practical on an end-to-end basis between SENET-DAXs, whether the ensuing bulk data transfer is to be accomplished as a Class I virtual connection, or as a Class II packet-switched call.

2.2.3.2 SENET-DAX Maximum Bulk Data Transfer Rate

The SENET-DAX concept has the capability to handle bulk data transfers at rates up to 1 million bits per second, provided adequate trunk capacity exists end-to-end between the originating and terminating SENET-DAXs. It should be noted however (see Table 2-2) that even with T1 carrier trunk groups, Class I

virtual connection of a 1 MB/S bulk data transfer would consume 64.77% of the frame capacity, while Class II packet-switching at this rate would consume 67.36% of the frame capacity. For the more frequent Bulk 1 transfers, these demands would persist for one second or less; for the relatively infrequent Bulk 2 transfers, these demands could persist for 1 to 100 seconds.

2.2.3.3 Practicality of Packet-Switching Bulk Data

It is considered practical to packet-switch bulk data; however, evaluation of the objectives as outlined in Question 3 of Section 2.2.2 (Objectives) indicates that there is not a single optimum answer to the question. As noted in Section 2.2.1, Bulk 1 data transfers are anticipated at rates of up to 16 per busy hour, while Bulk 2 data transfers are anticipated only once in ten busy hours. From Table 2-2, it can be noted that a minimal length Bulk 1 data transfer at the maximum rate anticipated (1 MB/S) would require approximately 2/3 of the total T1-based frame, but for only 1/100 second. Conversely, a maximum length Bulk 2 data transfer at the minimum rate anticipated (9600 B/S) would require only 2/3 of 1% of the frame capacity, but would persit for nearly three hours.

It is considered practical, and is recommended, that all Bulk 1 data transfers be accomplished as Class II packet-switched calls, regardless of the transmission rates involved, utilizing the deferability of Class II packets, and the SENET-DAX Class III throttling perogative as necessary during peak conditions. With respect to Bulk 2 data transfers, it is recommended that sensor data be Forward Error corrected and handled as Class I virtual connections at all rates up to 100 KB/s. For Bulk 2 data base transfers, Class II packet-switching is recommended at all rates up to 100 KB/S. Above 100 KB/S, and up to 1 MB/S, it is considered that Bulk data transfers are not practical with the SENET-DAX concept during peak

periods, unless the magnitudes involved do not exceed those previously defined (Section 2.2.1) for Bulk 1 data message lengths. To reduce both signaling and processing complexity, it is recommended that Bulk data transfers be constrained to 100 KB/S or less for SENET-DAX transmission.

2.2.3.4 Error Control Techniques for Bulk Data Transfers

Based on the evaluation, analysis, and recommendations given above, the control of errors in the transmission of Bulk 1 data, and in the transmission of Bulk 2 data bases will be accomplished utilizing the error control techniques described for Class II data packets in Section 3.2. Briefly, these techniques entail the SENET-DAX segmenting data messages (if necessary) into moderate-sized packets (typically 1000 to 2000 bits) and enveloping these packets based on the ADCCP technique. Each envelope includes individually self-identifying information, and each packet is individually protected by a cyclically-coded Frame Check Sequence (the Frame in the FCS being the packet, not the SENET master frame). These packet FCSs are utilized in an ARQ (Automatic Repeat Request) protocol requiring definite acknowledgment of receipt of valid packets. For Bulk 2 sensor data transfers, it is suggested that error control be based on Forward Error Correcting the data transfers, utilizing redundant information generated by each originating SENET-DAX, passed transparently (as a Class I virtual connection through each tandem SENET-DAX), and correcting as necessary at the terminating SENET-DAX.

2.2.3.5 Utilization of SENET-DAX Input Bulk Data Message Throttling

The SENET-DAX perogative of temporarily shutting off the input of Bulk data messages will be utilized first, in the establishment of the availability of the terminating subscriber for a bulk data transfer; and secondly, in the control of information transmission during peak traffic periods. In both cases, the intent is to control the total amount of memory required at the SENET-DAX for message storage.

2.2.3.6 Packet-Switched Class III Handling Re Class II Handling

It is suggested that Class III packet-switched data transfers be handled by consideration of their native precedence levels versus those of the Class II data message packets, and CCIS messages with which these types of bulk data packets are recommended to be grouped. While the inference might be drawn from the descriptions of the various traffic classes in Appendix A of the SOW that Class III data is of lower system precedence than the other classes, there is currently no indication of such a relationship from the subscriber standpoint. Similarly, for Bulk 2 sensor data transfers, it is suggested that these transmissions, recommended to be handled as Class I virtual connections, be treated at their native precedence levels versus those of the other Class I virtual connections (digitized voice, low-speed video, and FEC'D facsimile).

2.2.3.7 SENET-DAX Packet-Switched Bulk Data Accountability

Based on the recommendations above, SENET-DAX accountability for Bulk 1 data should be the same as for all other non-CCIS data (interactive, query/response, data base update, and narrative record traffic). The degrees of accountability suggested as required for the Class II data are indicated in Section 3.4; broadly, these requirements include complete journal and reference recording at both the originating subscriber's SENET-DAX, and, as a minimum, at the SENET-DAX on which the first addressee is homed (if not also at every SENET-DAX resposible for a delivery). Tandem SENET-DAXs would have only a momentary accountability for the individual packets of a Class II message, from acknowledgment of validated receipt from a SENET-DAX in the direction of the originator, until receipt of acknowledgment of validated reception of the packet by the next SENET-DAX toward the addressee. These degrees of accountability would also apply to Bulk 2 data transfers handled as Class II calls (suggested to be typically the transfer of an entire data base). Should the suggestion that Bulk 2 sensor data transfers be accomplished as Class I calls be accepted, the SENET-DAXs would

afford the calls the same degree of accountability as will be afforded any other Class I calls, as indicated in Section 3.1 (CCIS Procedures). Broadly, these procedures insure virtual connection accountability, but do not maintain accountability for the information transferred via the virtual connection. It has been suggested that sensor Bulk 2 data transferred in this Class I manner be provided with Forward Error Correction coding (as is suggested for Class I facsimile); however, no ARQ (Automatic Repeat Request) correction of errors in Class I Bulk 2 data transfers appears practical.

2.3 SERVICE DISTINCTIONS

2.3.1 Problem

To determine methods for locating the start and end of each class of traffic within a master frame as the boundaries between classes shift dynamically in response to changes in traffic intensity, traffic mix, etc.

2.3.2 Objectives

Any procedure for determining the boundary between classes must comply with the following objectives:

a. Operate independently of the location of the Class I/ Class II and Class II/Class III boundaries within each master frame (non-fixed boundaries).

b. Operate with negligible error. As long as the master frame remains synchronized, the location of the start of each class of traffic must be known in order to prevent the loss of data.

c. Operate with the minimum amount of channel overhead. The number of bits per master frame required to locate the start of each class of traffic should not be excessive.

2.3.3 Analysis and Results

2.3.3.1 Class Divisions

If it is assumed that each link is master frame synchronized, then the start of a Class I region is uniquely determined and is simply the first bit following the last bit of the coded sequence used to obtain frame synchronization.

A determination of the boundary between Class I and Class II traffic is simplified by the fact that sufficient information is contained within the master frame structure to locate this boundary

without using a special marker. In particular, the map of each Class I region which must exist at each DAX for each link terminating at that DAX provides the necessary information to determine each Class I/ Class II boundary. The Class II region could normally be defined to start with the first bit following the last bit of the last channel in the Class I region. As a check of the start of the Class II region, the first eight bits in this region must be the Flag sequence 01111110 which signals the start of a data packet per ADCCP. It should be noted that because of the self-identifying feature of Class II data packets, it would be possible for the DAX to detect their occurrence anywhere within a master frame. Although this capability may be exploited during an abnormal condition (e.g., during loss of master frame synchronization) there are no plans to transmit data packets in the Class I region during normal operation.

As a result of Section 2.2, the need for a Class II/Class III boundary is obviated since, as Section 2.2 proposes, all Class II and most of Class III traffic will be restricted to data packet form and will be handled as a single class of traffic, denoted simply as Class II traffic*. That portion of the Class III traffic which is not handled as Class II traffic will be handled as Class I circuit switched data. For further elaboration of this point, see Section 2.2.

The attractiveness of the chosen techniques for realizing the division between classes results from the following facts:

a. They are not directly dependent on special coded sequences and, consequently, are not directly affected by link error environment.

---

* All non-packet data terminals utilizing Class II service will have their data converted to packet form by the originating DAX prior to transmission through the network, and converted from packet form by the terminating DAX prior to transmission to the called terminal.

b. They depend on frame synchronization and the Class I map - two system functions which are maintained with extreme accuracy.

c. They can utilize DAX software and entail no additional channel inefficiency.

2.3.3.2 Efficiency Considerations

It is unlikely that in each master frame the packets queued for transmission will exactly fill the available Class II capacity. In those cases where the total bit requirement due to queued packets is less than the available Class II capacity, filler bits will be added to pad out the spare capacity. These filler bits will probably be ADCCP flag sequences or an all 1's bit pattern. In these situations where the queued packet bit requirements exceed the available Class II capacity and where an integral number of packets does not exactly fill the Class II region (which should almost always be the case), there will also be spare capacity. This spare capacity will always involve fewer bits than contained in the smallest packet queued for transmission. Note however that the spare capacity could involve a significant number of bits since in any master frame the smallest packet available for transmission could be hundreds of bits long. To avoid wasting this available capacity in situations where packets are queued for transmission, it is proposed that the next packet in the queue be split such that the first section of the split packet exactly utilizes the spare capacity available. subject to the constraint that the first section of a split packet never contains than the F, A and C fields. The remaining section would be transmitted in the next master frame at the beginning of the Class II region. If the subsequent master frame's Class II region is not large enough to contain the packet section remaining, this section could also be split in the same manner as the original packet. In fact, the splitting process could continue over many master frames, if necessary.

A description of control procedures required to handle this partial or split packet concept is provided in Section 3.2. It is recommended there that an ADCCP flag sequence serve as a Class I/Class II marker so that certain traffic conditions brought on by splitting packets can be handled efficiently. This flag sequence marker would be in addition to the flag sequence used to start the first unsplit packet in the Class II region. The use of a Class I/Class II marker does not impact significantly on the Class Division problem discussed in the previous sections since the marker to be used with splitting is the same bit pattern the DAX would see if splitting were not used. It should be noted though that when packet splitting is used, if the first packet in the Class II region is an unsplit packet then the first sixteen bits in this region will be two contiguous flag sequences

2.4 CCIS FIELDS AND FORMATS

2.4.1 Problem

The Class I region of a master frame is allocated to individual subscribers by assignments made via Common Channel Interswitch (CCIS) Messages. These CCIS messages are a sequence of short digital control messages used by the DAX's to communicate with each other in order to coordinate and control Class I switched traffic, efficiently, conveniently, and in real time. Toward this end, the various candidate CCIS message formats must be analyzed and the CCIS information fields defined. One possibility is to transmit the information in the form of packets, with fixed envelope and variable-length data fields, as suggested in the newly proposed Advanced Data Communication Control Procedure (ADCCP). ADCCP is a bit oriented control procedure where bits and bit patterns are used for control signals. Another approach would be utilization of character oriented procedures where the control signals are based on USASCII characters. This section will examine the suitability of using ADCCP to define the CCIS field sizes and format contents for interswitch transfer of control information between DAX's.

The choice of the ADCCP-based procedure is based in part on the desirability of utilizing the same procedure for CCIS messages, and for all Class II data messages, in order to minimize software design and processing demands. An even stronger reason for choosing ADCCP rather than USASCII procedures is the significantly greater

efficiency of the ADCCP. Roughly, at minimum practical lengths, ADCCP, with 5/6 ths of the bits required by USASCII (40 vs 48) can convey double the address information and double the control information (256 vs 128 in each case) with far greater error protection.

2.4.2 Objectives

The objective of this task is the efficient, convenient coordination and control of Class I traffic, in near real time. This requires achievement of the following:

1. A general mapping of the Class I region, which is efficient in terms of frame utilization and reasonable in terms of processing demand.

2. A definition of the CCIS message information fields and their sizes and the formats necessary for signalling between DAX's.

3. A determination of the link control signals that are to be used for routing, synchronization, error control, etc., and the sequence in which they are to be used to perform these functions.

2.4.3 Analysis and Results

2.4.3.1 CCIS Region

ADCCP-formatted CCIS messages are proposed to be handled identically to all other Class II data message packets. Consequently, CCIS messages are technically Class II messages; however, to ensure efficient, timely processing of Class I traffic, the CCIS messages may be handled at FLASH procedence level. Since the only two higher data precedences, CRITIC and ECP, can be expected to comprise less than one percent of the data traffic, the CCIS portion of the master frame will then directly follow the Class I Region and precede all other Class II data packets (see Figure 2-4.) CCIS messages contain information necessary to establish, terminate, and route Class I calls through the network, and to make changes in Class I allocations as a result of call terminations. CCIS messages are also used for other purposes such as Class II coordination, error control, synchronization, recovery, maintenance, and accountability, with respect to Class II data message traffic. In this section we will deal mainly with the fields and formats associated with Class I calls. It is not likely that the CCIS messages for Class II calls will be a subset of the messages we define for Class I operation. However, Class II CCIS messages, and those of synchronization and the like, will be the subject of Traffic Control Procedures (see Section 3.2).

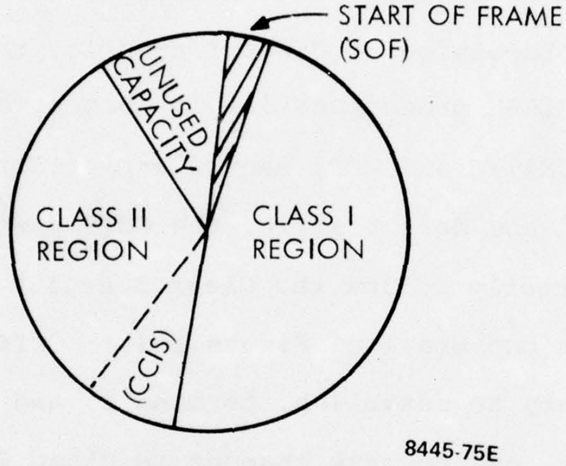

Figure 2-4. Master Frame Structure

The procedure to set up Class I calls involves initiating the messages to set up and confirm connections and then assigning specific portions of the Class I region to Class I subscribers. Since Class I calls will have durations on the order of minutes, several thousand master frames could result between the origination and termination of individual call connections. For example, assuming an average holding time of 5 minutes and a frame period of 10 msec, 30,000 master frames will result during the duration of a single call. Between the origination and termination of this call no information need be transmitted except that caused by realignment of the Class I region due to terminations of other calls. As calls are released the remaining calls are continually pushed up toward the beginning of the master frame to fill in the gaps as they occur. In the scheme a complete mapping of the total Class I contents frame by frame would result in significant transmission inefficiency. Therefore, in order to minimize this inefficiency, CCIS signaling will occur only when changes are required in the Class I region and only for those allocations affected, barring interruptions in transmission/reception. In the event of such an interruption, the first step taken by the DAX's will be to reestablish the maps of the Class I region maintained by the DAX's at each end of each interrupted link.

CCIS messages which are sent to describe changes to be made in Class I channel allocation are implemented in subsequent master frames, although it is desirable to make these changes at the earliest possible time. One reason for putting off the changes until later frames is the stringent demand on processing time required by connections and supervisory control made in the same frame interval in

which the CCIS messages are transmitted. Using the approach mentioned processing demand is considerably eased because connection time is extended by one or more frames per tandem line. However, the delay incurred may or may not be acceptable. A more important reason for changing the class I maps after receipt of the CCIS message is the necessity for changing both the transmit and receive maps with respect to exactly the same frame. Procedures for accomplishing this are described separately in Section 3.1.

2.4.3.2 Bit Oriented Versus Character Oriented Control Procedures

A data link between DAX's is shown in Figure 2-5 and provides for the movement of information between computers, terminals, or between computers & terminals. To accomplish this efficiently a control procedure is necessary to perform such functions as routing, synchronization, error control, and framing between stations on a link. Several link control procedures have been developed for this purpose; many of these standards are based on two separate control philosophies. One method uses control signals based on USASCII characters while the other uses bits & bit patterns for control. The control signals made up of the unique bit patterns bear little relationship to the USASCII code characters. The bit oriented control procedure (referred to here as ADCCP) permit improved features & characteristics relative to the USASCII character oriented procedure. However, a primary attraction in utilizing the ADCCP procedures and format is the greater efficiency for both CCIS and data message packet transmission which these procedures and format permit, relative to comparable USASCII prcoedures (see Section 2.1 - Framing & Enveloping).

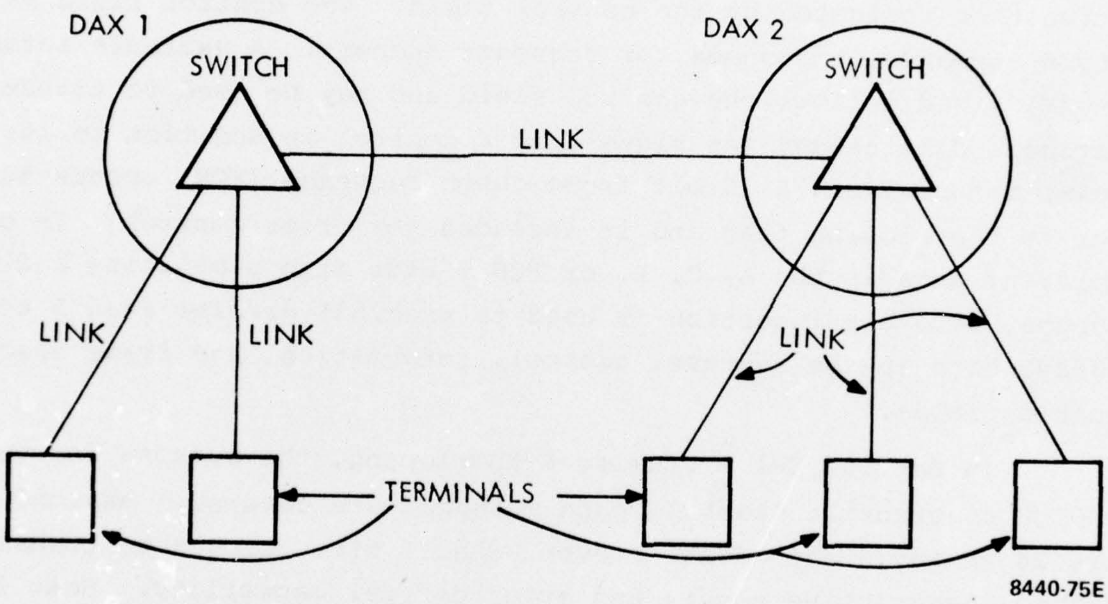

Figure 2-5. Portion of DAX Communication System

2.4.3.3 ADCCP Packet Structure

The frame structure of an ADCCP packet is shown in Figure 2-6. Also shown is the structure of a USASCII message block, the alternate approach we were considering for transmission of Class II data messages (including CCIS). Each ADCCP packet is preceded and terminated by an eight-bit flag which is the bit sequence 01111110. The flag (F) sequence which closes a packet may also be the opening flag sequence for the next packet. An address (A) field and control (C) field follow the opening flag. The addressed party accepts the packet and performs the function indicated by the control field. The control field may contain commands, responses, or sequence numbers. A variable information (I) field follows the control field and may be used to transmit additional data needed for signalling & control in addition to its regular information. A 16-bit frame check sequence (FCS) occurs just prior to the closing flag and is included for error control. In order to prevent data in the A, C, I, or FCS fields from simulating a flag sequence, zero bit insertion is used to prohibit greater than 5 consecutive bits in the address, control, information, and frame checking sequence fields.

In Section 2-1 - Framing & Enveloping, the reasons for our choice of ADCCP to transmit Class II data messages are covered. Basically, there is an advantage of ADCCP over USASCII with respect to transmission overhead, descriptive power, and error control capability. Some feeling for this can be seen from Table 2.3 where the different message fields for each approach are enumerated and the calculation for transmission overhead efficiency is shown for several cases.

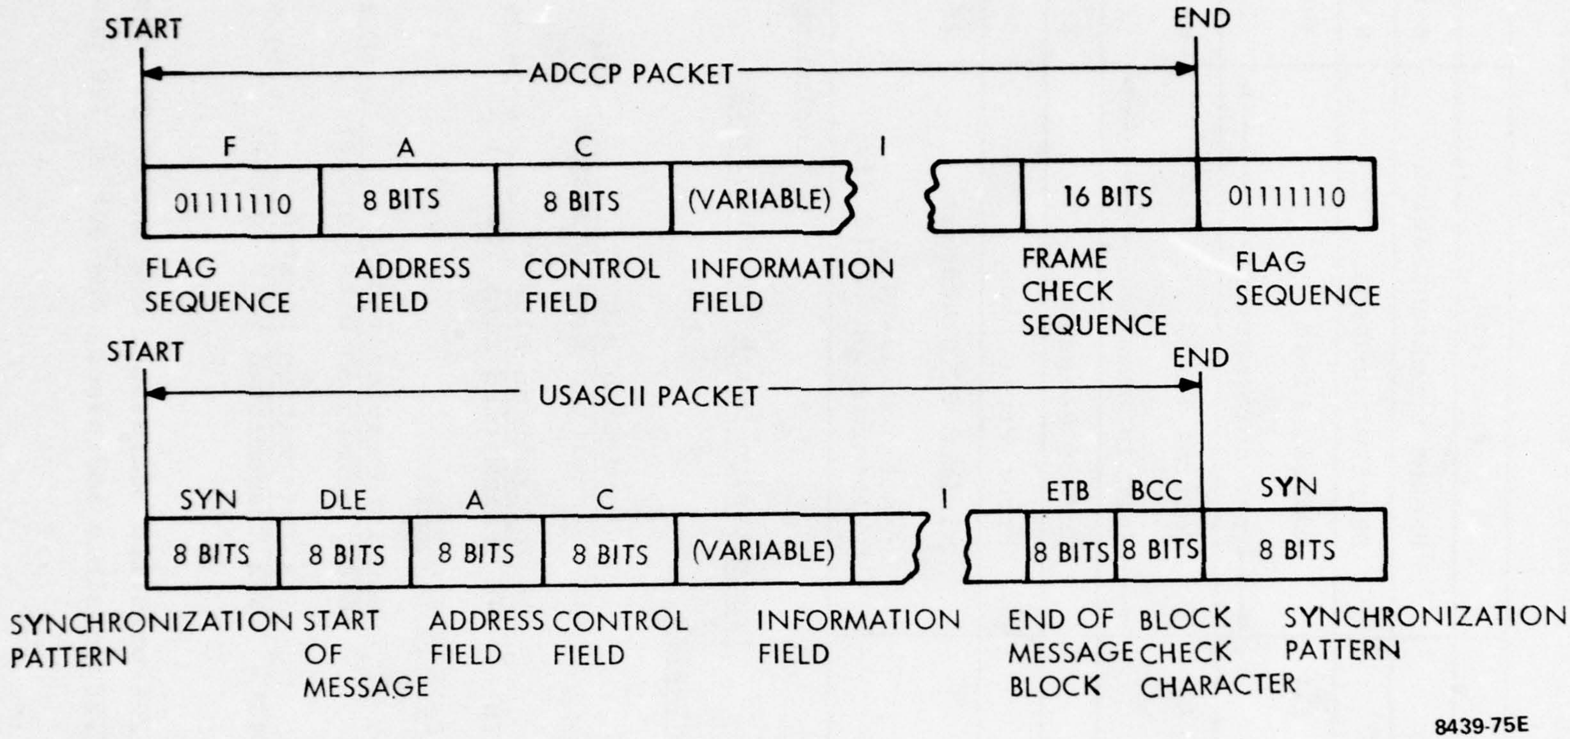

Figure 2-6. Basic Field Structure of ADCCP & USASCII Packets

TABLE 2-3. FORMAT STRUCTURE: ADCCP VS. USASCII

ADCCP	Function	USASCII
8-bit Flag (F)	Packet Synchronization	8-bit SYN*
-	Start of Message	8-bit DLE*
8-bit Address (A) (1 of 256)	Address Field ←(net information)→	8-bit Address (A)* (1 of 128)
8-bit Control (C) (1 of 256)	Control Field ←(net information)→	8-bit Control (C)* (1 of 128)
-	End of Message	8-bit ETB*
16-bit FCS	Error Protection	8-bit BCC*
40 bits	Minimum CCIS Message - not including Information Field (I)	48 bits

*including 1 parity bit

Transmission Overhead Efficiency

Minimum CCIS Message

$$ADCCP_{MIN} = \frac{\text{16 Information bits**}}{\text{40 Transmitted bits (gross)}} = 40\%$$

$$USASCII_{MIN} = \frac{\text{14 Information bits***}}{\text{48 Transmitted bits (gross)}} = 29.2\%$$

Information Field = 640 Bits

$$ADCCP_{640} = \frac{\text{640 Information bits**}}{\text{664 Transmitted bits (gross)}} = 96.4\%$$

$$USASCII_{640} = \frac{\text{638 Information bits***}}{\text{672 Transmitted bits (gross)}} = 94.9\%$$

**One out of 256 possible addresses, one out of 256 possible controls

***One out of 128 possible addresses, one out of 128 possible controls

Among the advantages of using ADCCP packets as CCIS messages include the following:

(a) Line control procedure can be defined in which the Information Field is code insensitive & transparent to the text being transmitted. Thus, the Information Field may contain any bit pattern & may be any number of bits long. Character-oriented procedure can be transparent & transmit any bit pattern; however, messages must be transmitted in multiples of 8 bits.

(b) Reliable error control in which control & data are protected by a 16-bit polynomial code. Generally a more powerful technique than orthogonal parity checking, the protection recommended in many character-oriented procedures.

(c) Efficiency in transmission overhead over that obtained using character-oriented procedure (see Section 2.1 - Framing & Enveloping)

(d) Flexibility in that additional link controls or responses can be added without changing the transmission format.

(e) Use of the same format structure as recommended for other class I messages, e.g., query/response, bulk data, etc.

2.4.3.4 Signaling Procedures*

The general approach toward signaling on the DAX network will be illustrated with the aid of Figures 4 and 5.** A call originates at switch 1 and is to be completed to a subscriber homing on switch 6. For illustrative purposes assume a commercial numbering plan although the principles presented here hold for military numbering plans as well. The subscribers dial anywhere from 7 to 10 digits of the form NAN - YYY - XXXX where NAN is the area code (A = 0, 1), YYY is the station #, and XXXX is the local subscriber #. At switch 1, the originating switch, the dialed digits are examined to see if (a) the call is to be completed to a foreign area code or (b) the call is to be completed to that area code but to a different switch within the home area.

In either case, the called subscriber is not to be found at that switch; however, each switch contains an area code table with all the stations in the network listed by area code. Therefore, if the called station is not in its locality the table will provide primary and alternate trunks by which to route the call to the next switch in the network in the direction of the called station. Here the same procedure is performed again. This is an example of spill forward

* This is a brief introduction to routing and signaling procedures. For a more detailed discussion, see Section 3.1-CCIS Procedures.

** This technique of routing on a switch-by-switch basis is applicable to a routing scheme which an access switch might use. Other schemes would be applicable for nodal switches in a backbone network.

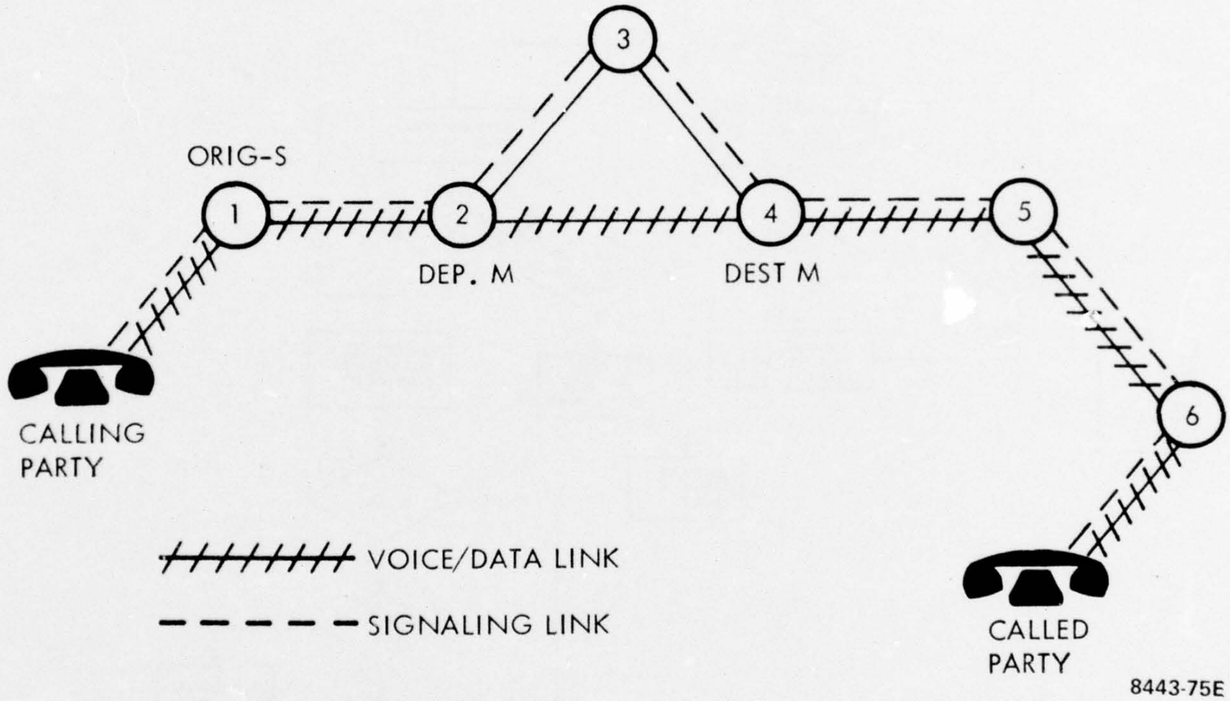

Figure 2-7. Portion of the DAX Network Showing Voice/Data and Signaling Links

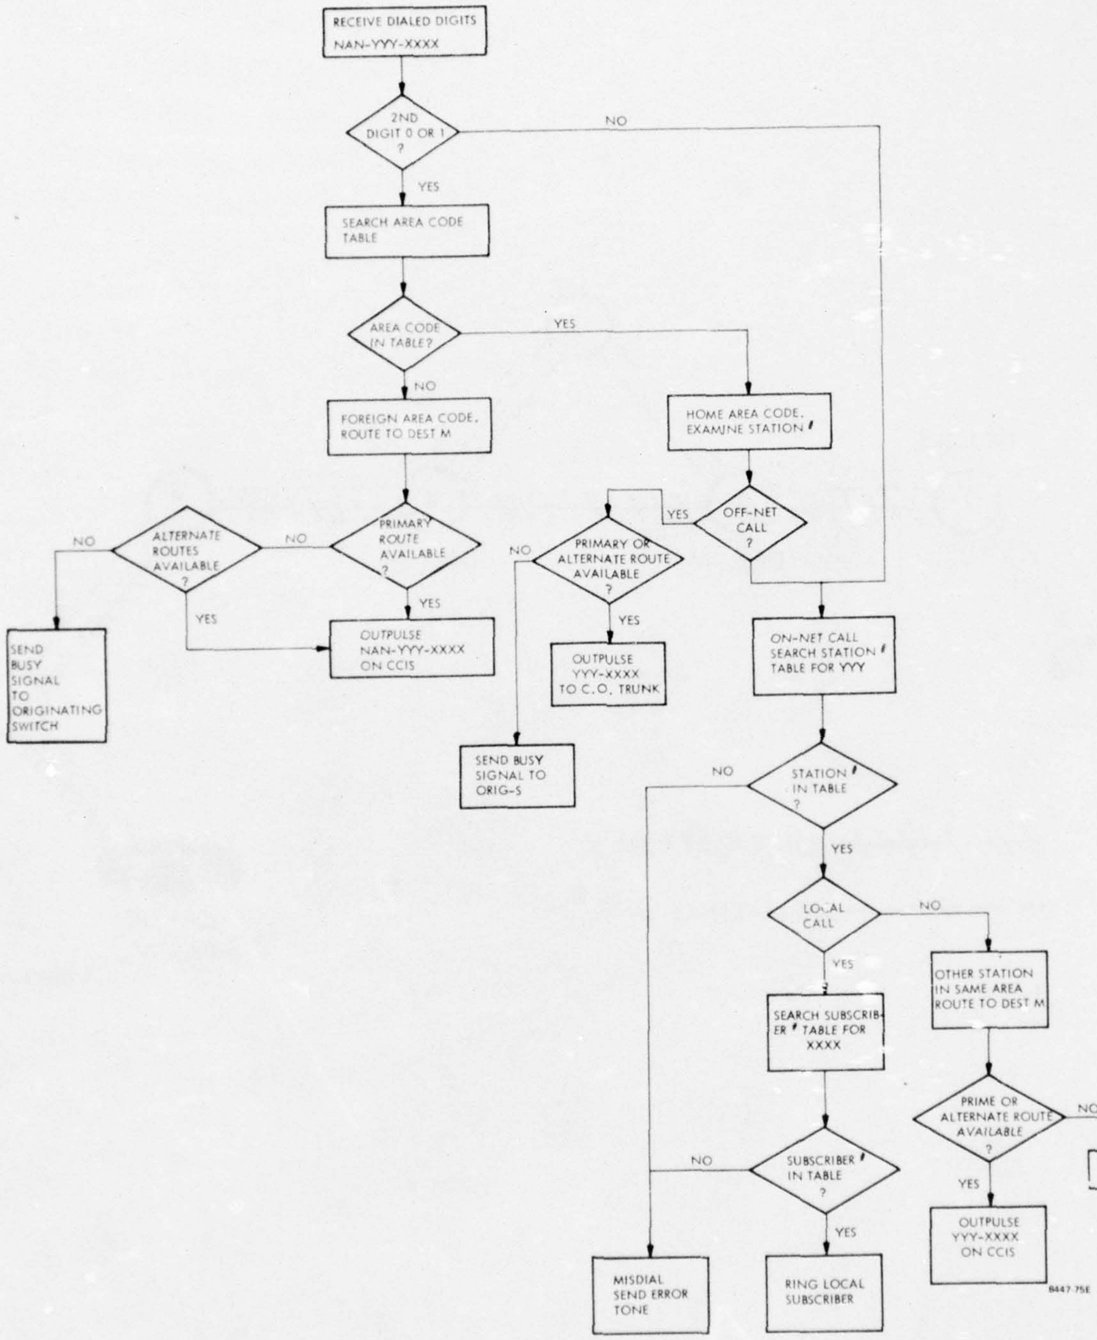

Figure 2-8. Translation and Routing Flow Chart for the DAX Network

routing control. The primary and alternate routing paths are chosen so that the call progresses through the network eventually reaching the terminal switch which completes the call. As links are established they are reserved until a complete path through the network to the called subscriber is determined.

In Figure 2-7 assume that the call setup has progressed one link from switch 1 to switch 2. The next link to be established is between switches 2 and 4. In the procedure, switch 2 is termed the point of departure of the message (DEPM) sent to establish the link while switch 4 is the destination of this message. Switch 1 remains the originating switch throughout the call setup procedure. In Figure 2-7, assume that the CCIS message has been blocked along the primary route from switch 2 to 4. Therefore, the call has been routed along a secondary route to switch 3 and from there to switch 4. However, when it came time to reserve the voice/data path, the primary route (1-2-4-5-6) was not blocked and was used to establish the path. This is an example of quasi-associate routing where the signaling path does not necessarily coincide with the voice/data path. For this type of routing, in addition to the content of each message the CCIS message must contain information to prevent the message from coming back on itself and/or getting lost within the network.

After the path between 2 and 4 is established, switch 4 becomes the new point of departure of the CCIS message (DEPM). The same procedure is used to establish the path between 4 and 5 and the call setup progresses in this fashion until the destination of message switch becomes the terminating switch. The terminating switch

then completes the call through to the called subscriber in the manner described in Section 3.1.

2.4.3.5 CCIS Message Types

Based on the signaling procedure described in the previous section, and the trunk signaling messages used in the AN/TTC-39, the following signaling messages are some of those necessary to establish and terminate Class I calls in the DAX network.

1. Call Initiate - used to reserve space for a channel in a trunk to the distant office and to initiate signaling to the called party

2. Call Complete - sent by the called party's switch to indicate that the called party is being rung. The message is sent to the 1st tandem switch and then progresses along from tandem switch to tandem switch until the switch which originated the call is reached. Also allows misrouting to be checked by repeating back the called party's # from terminating to originating switch for comparison

3. Call Answer - indicates that the called party has gone off-hook and that space should be allocated in the frame for the previously reserved channel

4. Release - specified trunk is released because one of the parties hung up

5. Preempt Release - specified trunk is released because the trunk has been pre-empted without reuse

6. Operator Recall - the switch at which recall occurs sends a recall message to the originating or terminating switch to recall the operator to a call in progress

7. Acknowledge - acknowledges the correct receipt of a signaling message from the sending switch. Message is identified by the trunk and sequence number

8. Non-Acknowledge - the message as received contains errors

9. Glare - signifies that a request to seize the same trunk has been made by two switches at the same time

10. Out-of-Service - specific trunk is marked out of service for maintenance purposes

11. All Trunks Busy, Equipment Busy, Invalid Route - originated by a tandem switch in the network in response to a Call Initiate Message if the call is blocked for some reason in the network. The messages are sent back to the originating switch. Tandem switches along the route cancel path reservations; the originating switch notifies its call originating subscriber

12. Unassigned or Invalid Number, Incompatible Connection, Called Party Unavailable – same as for (11) only refers to the call being incomplete due to line/subscriber problems and not network problems.

13. Call Forwarding – informs home switch that the local subscriber wants his calls automatically transferred to another number. The switch directs his call to the new number

14. Acknowledge Call Initiate – acknowledges that a specific trunk at the distant switch has been seized and a path reserved in the link map

15. Synchronization Messages – used to synchronize sending and receiving sides of the transmission link

16. CCIS Keep Alive – used to report that no changes have occurred in the Class I region during the previous interval (T frames)

17. Maintenance – messages used to report failures in power supply, processor, trunk carrier, terminal, etc.

18. Traffic Statistics – provides traffic data on trunk seizures, completed calls, blocking, misdials, and other traffic information

19. Realignment – realignment of the Class I region to fill in missing gaps due to call terminations

20. Reinitiate Synchronization – initiate the synchronization procedure when out-of-synchronization is detected.

2.4.3.6 CCIS Message Fields and Format

The CCIS messages will be sent in the form of packets according to the structure specified in the ADCCP standard. This structure, shown in Figure 2-6, consists of a flag (the eight bit pattern 01111110) followed by an address and control field. Information data follows the control field and a 16 bit redundancy check always follows the text (if any). Figure 2-9 indicates the message format and their field sizes for typical signaling messages.

Consider the format for the Call Initiate Message. The message fields are described as follows:

1. Flag – bit sequence 01111110

2. Address Field – broken up into a 7 bit destination of message for which the message is bound and a 1 bit flag which signifies whether this is a CCIS message or a Class II data packet.

3. Control Field – uses information transfer format which will contain transmitting and receiving CCIS message sequence numbers

4. Message Type – up to 256 message types allowable

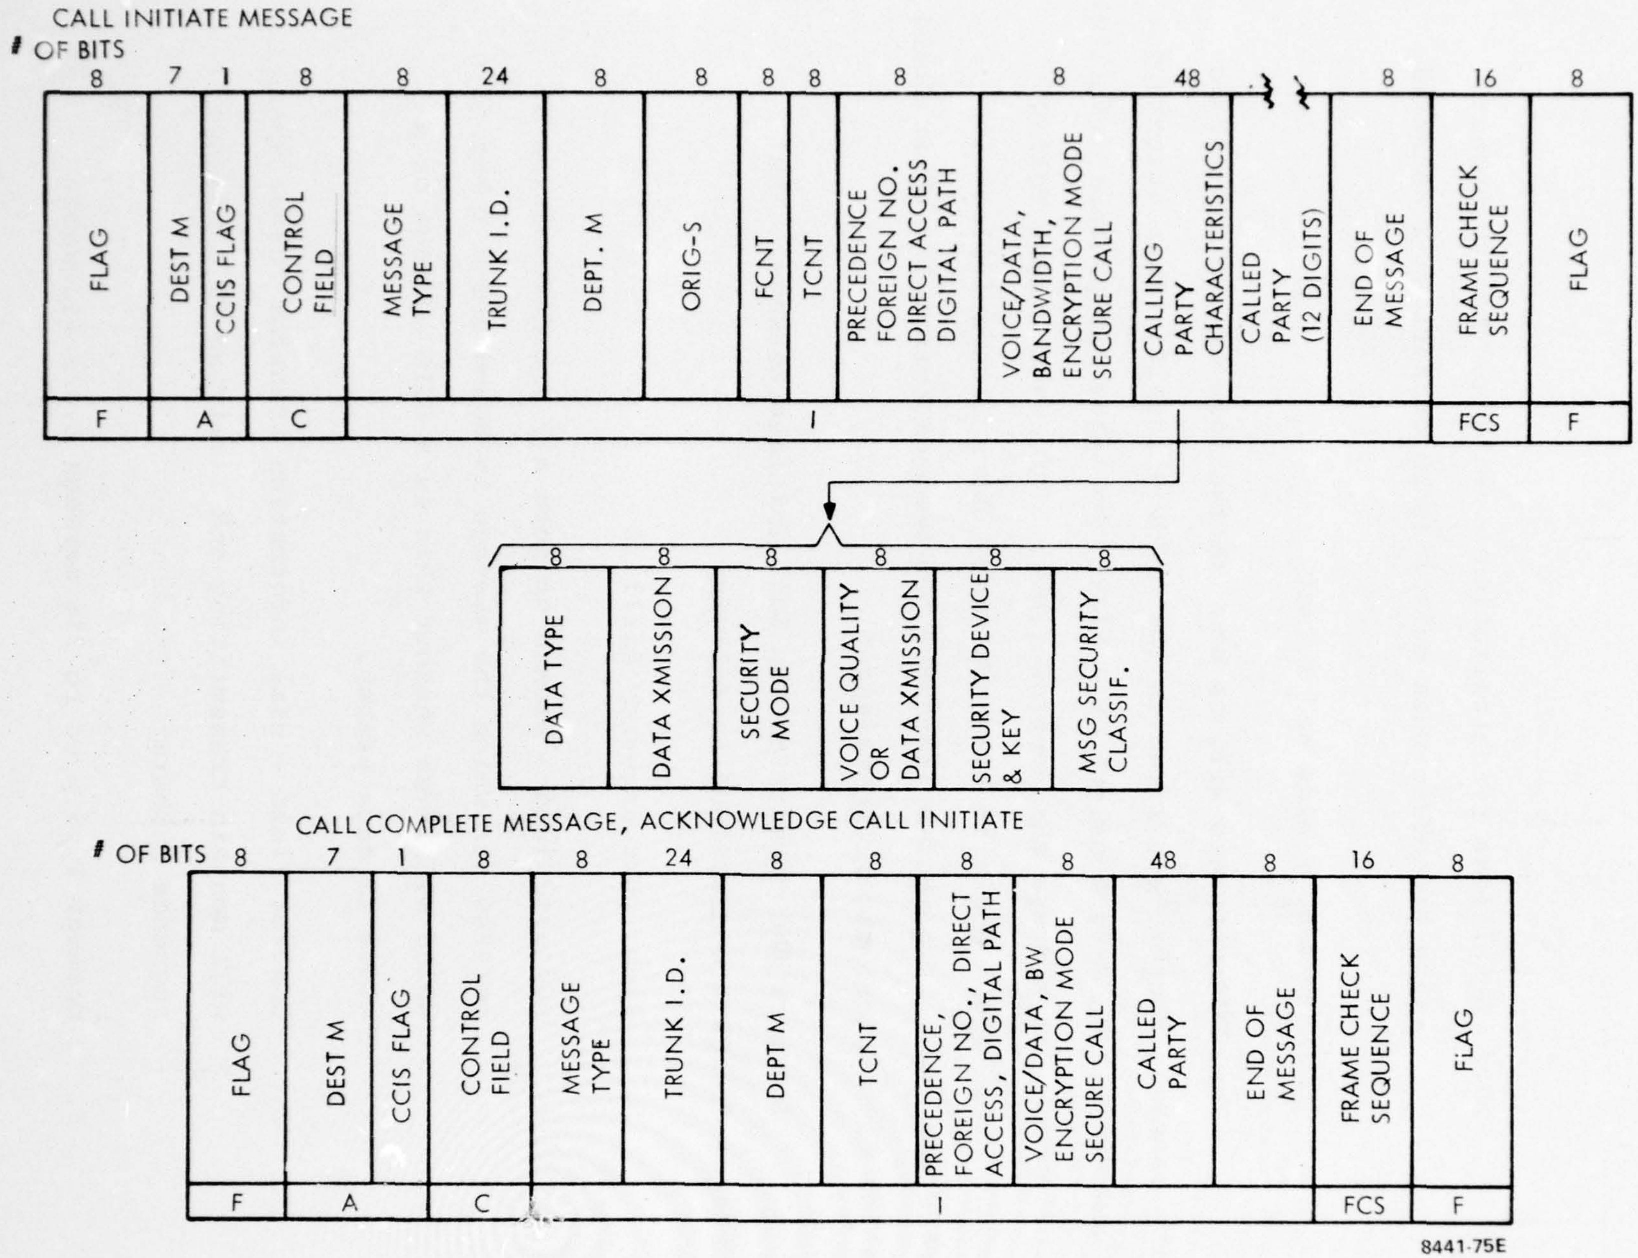

Figure 2-9. CCIS Formats and Field Sizes

CALL ANSWER, RELEASE, PREEMPT RELEASE, OPERATOR RECALL

# OF BITS	8	7	1	8	8	24	8	8	8	8	8	16	8
	FLAG	DEST M	CCIS FLAG	CONTROL FIELD	MESSAGE TYPE	TRUNK I.D.	DEPT M	ORIG S	TCNT	PRECEDENCE	END OF MESSAGE	FCS	FLAG
	F	A		C	I							FCS	F

ACKNOWLEDGE, GLARE, OUT OF SERVICE, SYNC, CALLED PARTY UNAVAILABLE, INVALID #, INCOMPATIBLE CONNECTION

# OF BITS	8	7	1	8	8	24	8	8	8	16	8
	FLAG	DEST M	CCIS	CONTROL FIELD	MESSAGE TYPE	TRUNK I.D.	DEPT M	ORIG S	TCNT	FCS	FLAG
	F	A		C	I					FCS	F

ALL TRUNKS BUSY, EQUIPMENT BUSY, INVALID ROUTE

# OF BITS	8	7	1	8	8	24	8	8	8	16	8
	FLAG	DEST M	CCIS	CONTROL FIELD	MESSAGE TYPE	TRUNK I.D.	DEPT M	ORIG S	TCNT	FCS	FLAG
	F	A		C	I					FCS	F

8438-75E

Figure 2-9. Cont.

5. Trunk I.D. – identifying number between the two communicating switches which specifies the starting bit of the Class I channel and its width

6. Dept M – identifies the sender of the message

7. Orig S. – identifies the originating switch

8. FCNT – counts the times a busy indication is given as the call progresses through the network toward the final destination

9. TCNT – counts the switches each message goes through to establish a voice/data path

10. Precedence – 3 bits representing the level of precedence, one bit for whether called party # uses DAX numbering plan or belongs to a commercial or foreign network, one bit for indicating a direct access call and one bit indicating whether the # contains 7 or 10 digits

11. Binary Indication of Voice or Data, Bandwidth, Mode of Encryption and Secure Call

12. Data Type - indicates line type mode format, code, and the # of stop bits for asynchronous codes

13. Data Transmission Type - defines type of modulation, data rate, etc.

14. Security Mode - one character indicates one of 3 modes for all calls, end-to-end, link-by-link, or non-secure

15. Voice Quality - estimates voice quality of a call based on the # of security devices a call traverses. When a threshold is exceeded, the call is terminated

16. Security Device & Key - security device (HY-11, KY-3) or key variable associated with a call

17. Message Security Classification - defines the security classification for the call

18. Called Party - up to 12 characters are permitted to allow for local or foreign numbers (commercial, NATO, etc.)

19. End of Message - control character EOT - end of transmission

20. Frame Check Sequence - 16 bit redundancy check.

The format and field sizes for other CCIS signaling messages are indicated in Figure 2-9. In general, overall CCIS message length is a function of message type and the manner each message is used in the CCIS procedures.

One modification of this format is to place the End of Message field close to the beginning of the message rather than at the end just prior to the Frame Check Sequence. In this case, not only would the field have to contain a delimiter signifying end of transmission, but, also, an indication of the number of total bits in the message. An advantage to this change in format structure is that a complete message up till the Frame Check Sequence can be read into memory quickly and checked for errors before any action is taken on other bit fields. With the format indicated in Figure 2-9, software would have to search bit by bit until the end of message field is identified.

2.4.3.7 *CCIS Procedures*

Section 3.1 will involve incorporation of this material into the development of CCIS procedures to establish and support Class I service. As a result of the future impact of changes in other factors, such as routing or preemption, possible redefinition in the sizing of the CCIS fields or in the format content could result.

CCIS messages for Class II coordination, synchronization, error control, maintenance, recovery and other message types, will be analyzed for format and content in Section 3.2.) New message types, which may only become evident as the DAX system is specified in greater detail, should be handled in the same fashion.

SECTION 3

ANALYSIS OF PROCESSING REQUIREMENTS

SECTION 3
ANALYSIS OF PROCESSING REQUIREMENTS

3.1 CCIS PROCEDURES

3.1.1 Problem

The Class I region of a master frame is allocated to the switched traffic of individual subscribers. The control and coordination of these individual subscriber "channels" is accomplished through the use of Common Channel Interoffice Signaling (CCIS) messages. The various individual CCIS message fields and formats are discussed in Section 2.4. The problem addressed in this section is the determination of the CCIS message sequences required to accomplish the various control and coordination functions.

3.1.2 Objectives

a. To develop a set of CCIS procedures which will satisfy the following criteria;

   1. Permit processor software determination of the varying Class I/Class II boundary on a frame-by-frame basis.
   2. Permit dynamic remapping of the switched data channels in the Class I region on a DAX-to-DAX interactive basis.
   3. Provide for the establishment and breakdown of switched data calls of routine precedence.
   4. Provide for handling of pre-emptive switched data calls.
   5. Provide for DAX-to-DAX transfer of administrative traffic, e.g. Automatic Message Accounting (AMA) data, SYSCON messages, etc.
   6. Minimize cross-network time from dialing of the last digit by the subscriber to provision of ringback to the subscriber.
   7. Maximize utilization of Class I region; i.e., allocation of channels no larger and no earlier than necessary, dropping of channels as soon as possible after release of call(s), repacking of Class I region as soon as channel(s) are dropped.

3.1.3 Approach

In the course of establishing the CCIS procedures, the need for trade-offs became apparent in the areas of signaling and routing control and frame allocation; these are discussed in succeeding paragraphs.

3.1.3.1 Originating Office vs. Spill Forward Control

The methods considered are described below:

3.1.3.1.1 Originating Office Control

Pure originating office control implies that the originating switch controls the signaling to each tandem switch (and terminating switch) as the call progresses through the network, with each tandem switch attempting to route the call over a primary route. In the event that blocking is encountered on one of the primary routes, the originating switch will then reinitiate the call over one of its alternate routes. Thus, in a conventional network, each tandem switch can cut through the speech path and release its call processing register as soon as an outgoing link is established from the switch. The originating switch, on the other hand, must keep the call processing register dedicated to the call until the called subscriber's terminal is located at the terminating switch. One principal advantage of originating office control is that the signaling to any tandem switch need include only sufficient digits for that switch to translate and route the call. In a network with inband signaling, that tandem switch would then connect the voice path through to the next switch in the route and become transparent to subsequent signaling. In a CCIS network however, there is no transparency because the signaling is done on a link-by-link basis and the number of CCIS messages required to set up a call becomes unacceptably high as the number of tandem links increases. Another advantage of originating office control is that the alternate routing capability is greater than in a spill-forward network, particularly when the originating office has a high degree of connectivity to other nodes because if blocking is encountered the call can be backed up toward the originating switch and alternate routed from there. This advantage holds true for both signaling methods, inband and CCIS.

AD-A039 548 GTE SYLVANIA INC NEEDHAM HEIGHTS MASS ELECTRONIC SYS--ETC F/G 17/2

SENET-DAX STUDY. VOLUME 1.(U)

JUN 76 DCA100-75-C-0071

UNCLASSIFIED FR76-1-VOL-1 NL

2 OF 3

AD A039548

3.1.3.1.2 Spill Forward Control

Spill forward control implies that the complete control of the call, both signaling and routing, is spilled forward from one switch to the next as the call progresses through the network. In a conventional (analog) system this presents two advantages over originating office control:

a. The signaling is reinitiated at each switch and is therefore not subject to the same degree of attenuation and distortion as in originating office control, particularly on calls passing through many links.

b. The call processing register (software) at the originating switch can be released as soon as the control is spilled forward to the next switch, resulting in less demand on processor resources, both in memory requirements and processor loading.

3.1.3.1.3 Suggested Approach

The approach that is suggested for the DAX is a modified originating office control with spill forward signaling. In this method, any switch will forward digits (signal) only to an adjacent switch. Each tandem switch will attempt to forward the call over alternate links if the primary link is blocked. If no idle or pre-emptible path is available out of the switch, both routing and signaling control of the call will revert back to the preceding switch. If no alternate or outgoing paths are available out of that switch, the call may either be dropped or revert back to the next preceding switch, etc. Whether or not control will be allowed to pass backward through more than one switch has not yet been decided, but represents a trade-off between increased alternate routing capability and increased CCIS message requirements and call processing register holding time.

3.1.3.2 *Class I Region Channel Allocation*

A decision is required in the process of establishing a call, as to when to insert in the Class I Region the channel that will eventually be required for the call. Several alternatives have been considered, each having its own advantages and disadvantages.

a. Progressive Forward Channel Allocation
b. Delayed Forward Channel Allocation
c. Progressive Backward Channel Allocation
d. Delayed Backward Channel Allocation

3.1.3.2.1 Progressive Forward Channel Allocation

In this method, the channel for the call is allocated on a link-by-link basis as part of the signaling spill forward process. Differential delays can be introduced in both the sending and receiving DAX so that the actual allocation in the Class I regions of the link in both directions (DAX 1 to DAX 2 and DAX 2 to DAX 1) will be synchronized to within one frame period.

The advantages of this method are:

a. Of the four alternatives considered, it is the simplest to implement.
b. It insures that the speech path will have been cut through before the start of subscriber ringing so that the calling subscriber will hear ringback at the earliest possible time.

Disadvantages of the method are:

a. It does not maximize Class I region utilization because the tandem links are loaded with the channel before it is known that the call can be completed.
b. In addition, if allocation has been made and the call cannot be completed, the Class I region must be re-compacted when the unused allocation is dropped.

3.1.3.2.2 Delayed Forward Channel Allocation

In this method the channel for the call is not allocated until the sending switch receives a message from the receiving switch that it can either terminate or extend the call.

The advantages of this method are:

a. It insures that the speech will have been cut through before the start of subscriber ringing so that the calling subscriber will hear ringback at the earliest possible time.

b. The allocation portion of the procedure can be overlapped with the process of forwarding the call to the next switch so that minimal extra cross-office delay is incurred at tandem switches.

c. The required frame capacity may be used for Class II data until it is subsequently allocated to the Class I channel. (In the order of 100 frames in the case of a satellite link.)

d. It provides a one-link look ahead capacity to detect call blocking before allocation of the channel.

The disadvantages of the method are:

a. It is more difficult to implement than the method described in 3.1.3.2.1, although if blocked calls are permitted to back up no more than one switch, the additional complexity is not too great.

b. It does not maximize Class I region utilization because the tandem links are loaded with the channel before it is known that the call can be completed.

c. In addition, if allocation has been made and the call cannot be completed, the Class I region must be re-compacted when the unused allocation is dropped.

3.1.3.2.3 Progressive Backward Channel Allocation

In this method the channels are reserved as described in 3.1.3.2.2, but no allocations are made until the call reaches the terminating switch, at which time the channels are allocated in both directions a link at a time tracing backward along the path of the call.

The advantages of this method are:

a. The required frame capacity may be applied to the Class II region until it is subsequently allocated to the Class I channel.

b. A channel will not be allocated in any link until it has been established that the called subscriber can be rung.

The disadvantages of the method are:

a. It introduces an additional network delay (≈100 frames for satellite links) from time of dialing last digit to hearing ringback.

3.1.3.2.4 Delayed Backward Channel Allocation

In this method the channels are reserved as in 3.1.3.2.3 but no allocations are made until the called party answers the call. Instead, when the call reaches the terminating switch the switch simultaneously starts ring forward to the called subscriber and transmits a quasi-associated message backwards towards the originating switch informing it that the called party is being rung. The originating switch then connects the calling party to a ringback tone.

When the terminating switch detects answer supervision from the called subscriber, it then initiates the backward allocation sequence as in 3.1.3.2.3.

This method has the virtue of not tieing up the Class I region while the calling subscriber is being rung. For example, in a call which rings for 12 seconds before being answered, the delayed backward allocation method would save 1200 frame intervals of channel allocation per link over the method of 3.1.3.2.3.

It is suggested that in each link the channel allocations be symmetrical in both directions. Although it is theoretically possible to delay the allocations in the forward direction until the backward allocations are accomplished end to end, it is felt that this would complicate the maintenance of the frame maps tremendously as well as opening up the possibilities or such phenomena as unidirectional link blocking or pre-emption. Also, the effect of end-to-end encryption has not yet been addressed. This too might serve to make differential allocations inadvisable.

3.1.3.2.5 Suggested Approach

It is proposed to incorporate both the methods described in 3.1.3.2.3 and 3.1.3.2.4 to be determined on a per call basis at the terminating switch of each call as a function of terminal classmark. Where advanced subscriber equipment is available from which both ring supervision and answer supervision signals are available to the switch, the delayed backward allocation method of 3.1.3.2.4 will be used; otherwise the method of 3.1.3.2.3 will be employed.

3.1.4 Progress

This paragraph describes the CCIS procedures and message sequences currently proposed for establishing and breaking down Class I calls.

3.1.4.1 Establishing Class I Calls

Class I calls will be established in two distinct phases. The call set-up phase, described below, involves establishing a route through the network to the called subscriber, ringing the called subscriber, and returning ringback to the calling subscriber. The channel allocation phase, described in 3.1.4.1.2, involves the allocation of channels (i.e. establishing a bi-directional speech path) between the two subscribers. The latter phase is normally delayed until the called subscriber goes off-hook (i.e. answers the call). If the called subscriber's equipment does not have the capability of returning an answer signal to the terminating DAX then the channel allocation phase will not be delayed but rather, initiated concurrently with ringing the called subscriber. The ringback will originate at the called subscriber's equipment and will be propagated backwards towards the calling subscriber over the newly allocated channels. The paragraphs below consider a Class I call with delayed allocation.

3.1.4.1.1 Call Set-Up Phase

Figure 3-1a depicts the overall CCIS signaling sequence used to set up a Class I call through a hypothetical sub-network. The sub-network connectivity and link availability and routing strategy have been established so as to provide examples of blocking along the

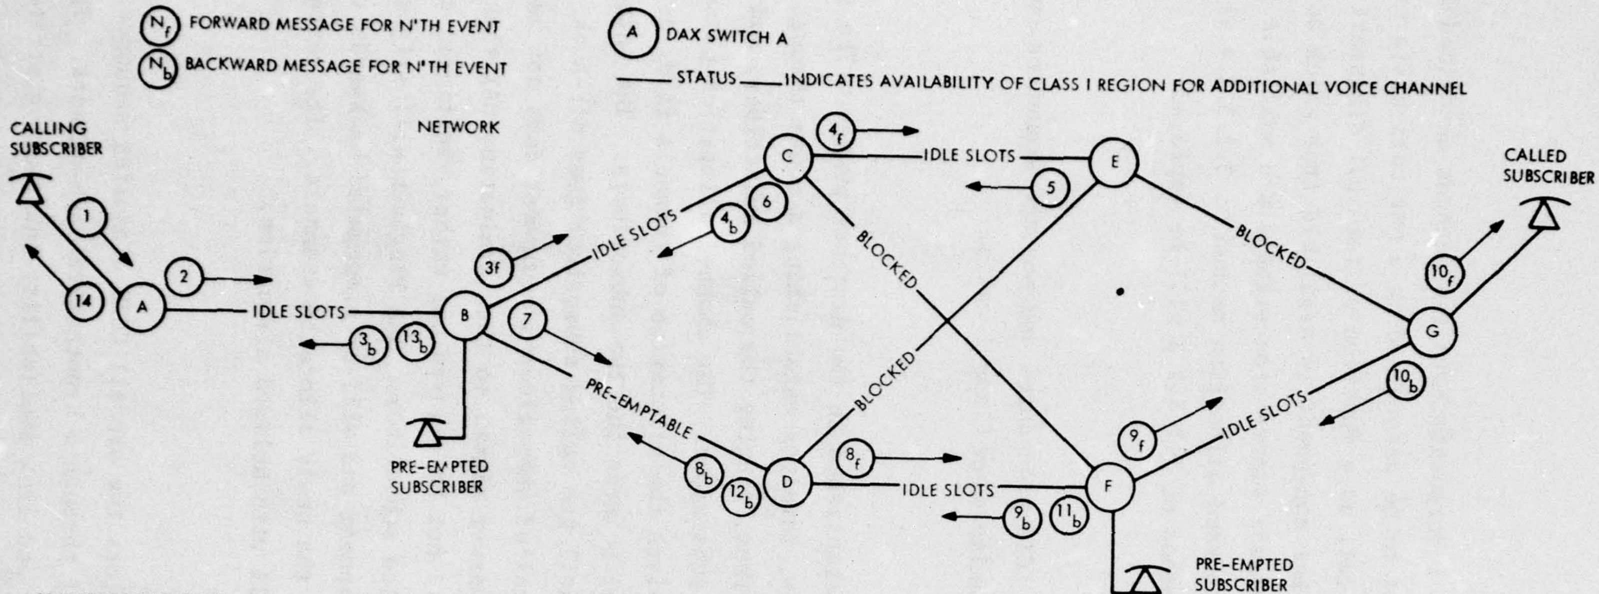

SEQUENCE OF EVENTS

1. CALLING SUBSCRIBER DIALS LAST DIGIT
2. RESERVE A-B CHANNEL
3. f - RESERVE B-C CHANNEL/b - A-B CHANNEL RESERVED
4. f - RESERVE C-E CHANNEL/b - B-C CHANNEL RESERVED
5. BLOCKING ENCOUNTERED, C-E CHANNEL DENIED, REROUTE
6. BLOCKING ENCOUNTERED, UNRESERVE B-C CHANNEL, REROUTE
7. RESERVE PREEMPTIVE B-D CHANNEL
8. f - RESERVE D-F CHANNEL/b - B-D PRE-EMPTIVE CHANNEL RESERVED
9. f - RESERVE F-G CHANNEL/b - D-F CHANNEL RESERVED
10. f - RING FORWARD TO SUBSCRIBER

(10) b TO (13) b QUASI-ASSOCIATED CCIS RINGBACK MESSAGE FROM DAX G TO DAX A

(14) RINGBACK SENT TO CALLING SUBSCRIBER FROM DAX A

9111-75E

Figure 3-1a. CCIS Procedures for Call Set-Up Phase

primary route and link pre-emption along the secondary route. The sequence of events is shown on Figure 3-1a but a few words here are in order.

One basic ground rule of the proposed procedures is that a DAX switch will never modify the Class I region of a link until it has first notified the adjacent DAX of its intent to do so and has received a confirming response from the adjacent DAX via CCIS and has acknowledged the response via CCIS. This positive handshaking approach is taken in order to prevent an error in any CCIS message from causing perturbations (and therefore noise) in all portions of the Class I region affected by the CCIS message. For example, pre-empting or dropping the first channel can affect the starting bit position of every other channel in the Class I region of that frame. This approach results in a request, a response and an acknowledge CCIS message in every tandem DAX to DAX transaction, except at the terminating DAX where the first backward allocation request serves as the response to the last forward reservation request. In general there are two transactions between every pair of DAX in the route. The acknowledge messages are not shown in Figure 3-1a because they would add more confusion than clarification. They do show up however in Figure 3-1b, about which more will be said.

The sequence of events described in Figure 3-1a is:

1. DAX A accumulates the dialed digits, translates and finds the primary route is to DAX B.
2. DAX A signals DAX B its intent to reserve a channel in the A-B link, and waits for a response.
3. DAX B translates the dialed digits, finds primary route is to DAX C. DAX B signals forward to DAX C its intent to reserve a channel in the B-C link and signals backward to DAX A its agreement to reserve the channel in the A-B link. DAX A then acknowledges DAX B's response and both DAXs reserve the channel simultaneously in a subsequent frame.
4. In the meantime DAX C translates the dialed digits, finds the primary route is to DAX E, signals forward and waits for a response from DAX E.

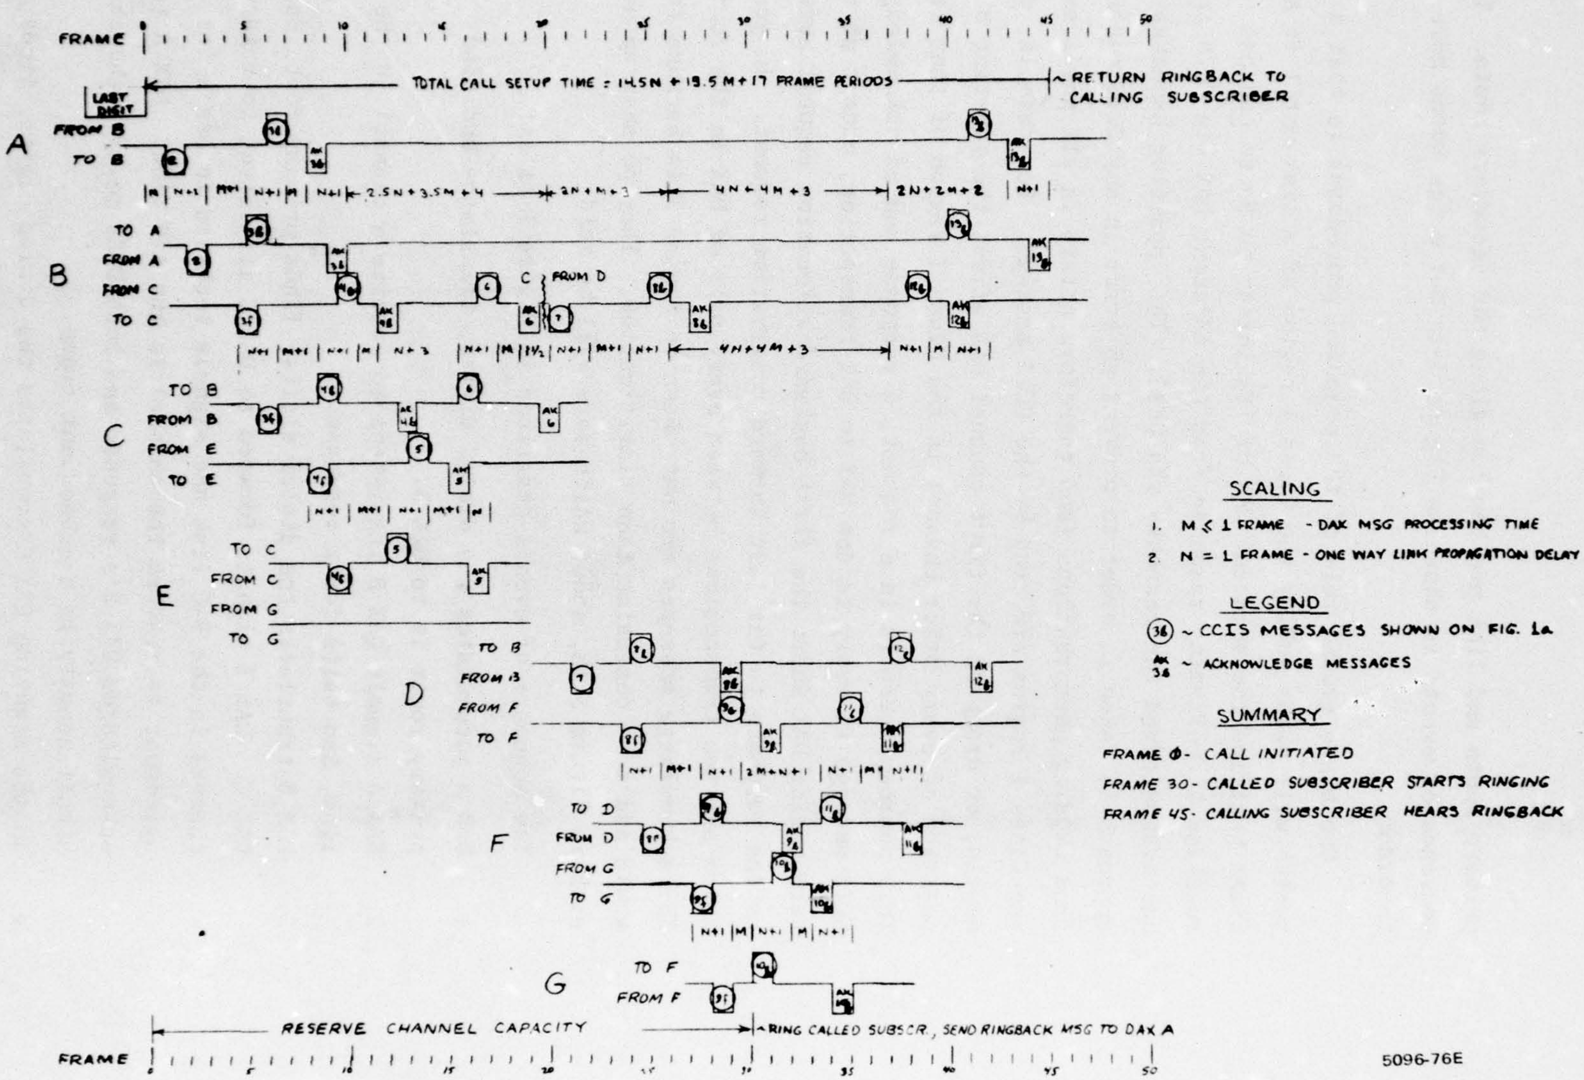

Figure 3-1b. CCIS Timing for Call Set-Up

5. DAX E translates the dialed digits, finds both the primary and alternate routes blocked and signals backward to DAX C that the call is blocked, that it (DAX-E) won't reserve the C-E channel and that DAX C should reroute the call.
6. DAX C receives the blocking message, attempts to reroute and finds the alternate route to DAX F is blocked. DAX C then signals backward to DAX B that the call is blocked, that the B-C channel reservation should be canceled and that DAX B should reroute the call.
7. DAX B receives the blocking message, finds an alternate route to DAX D requiring the pre-emption of a channel containing an existing call, and signals forward to DAX D its intent to reserve the channel in the B-D link.
8. DAX D receives the reservation message, translates the dialed digits, finds a primary route to DAX F, signals forward to DAX F its intent to reserve a D-F channel, and signals backward to DAX B its agreement to reserve the channel in the B-D link.
9. DAX F translates the dialed digits, finds a route to DAX G, signals forward to DAX F and backwards to DAX D.
10. DAX G translates the dialed digits and finds that it can terminate the call. Since the called subscriber's equipment is capable of sending an answer supervision signal back to DAX G, DAX G then sets up to ring the called subscriber and sends a quasi-associated message to DAX A, via DAX F, D and B, to return ringback to the calling subscriber.

Figure 3-1a depicts the timing of the message sequences described above. Examination of 3-1b reveals that the total cross-network time to establish a call does not equal the sum of the signaling sequences across individual links but is significantly shorter due to overlapping sequences. For the call described in Figure 3-1, the sum of the individual sequence times would be 26N + 26M + 33.5 = 88.5 Frame Periods whereas the actual time to establish the call is 14.5N + 13.5M + 17 = 45 Frame Periods or 450 MSEC, where:

N = one way propagation delay across a link = 1

M = individual DAX processing time = 1.

3.1.4.1.2 Channel Allocation Phase

Figure 3-2 depicts the overall CCIS signaling sequence used to allocate channels to the call set-up in the preceding paragraph. Again, the acknowledge messages are not shown in order to avoid confusion.

1. DAX G receives an answer supervision signal from the called subscriber.
2. DAX G initiates the backward allocation sequence by signaling backward to DAX F its intent to allocate the F-G channel which DAX F had requested it to reserve.
3. DAX F receives the allocate message and signals backward to DAX D to allocate the previously reserved D-F channel. DAX F then signals DAX G its intent to allocate the F-G channel.
4. DAX D receives the allocation message, signals backward to DAX B to pre-empt and reallocate the previously reserved for pre-emption B-D channel and signals forward to DAX F its agreement to allocate the D-F channel. In addition, DAX D examines the status of the pre-empted call, finds that the pre-empted call utilized another channel in the D-F link and signals DAX F to pre-empt that channel. DAX F drops the channel and sends the pre-empt tone to the pre-empted subscriber (who is local to DAX F).
5. DAX B receives the pre-empt and re-allocation message from DAX D, signals backward to DAX A to allocate the previously reserved A-B channel, and signals backward to DAX D its agreement to pre-empt and reallocate the B-D channel. In addition, DAX B examines the pre-empted channel, finds that the pre-empted call terminates locally and sends pre-empt tone to the pre-empted subscriber.
6. DAX A receives the allocate message from DAX B and signals DAX B its agreement to allocate the previously reserved A-B channel, removes the ringback signal and connects through to the calling subscriber, enabling two-way conversation to ensue.

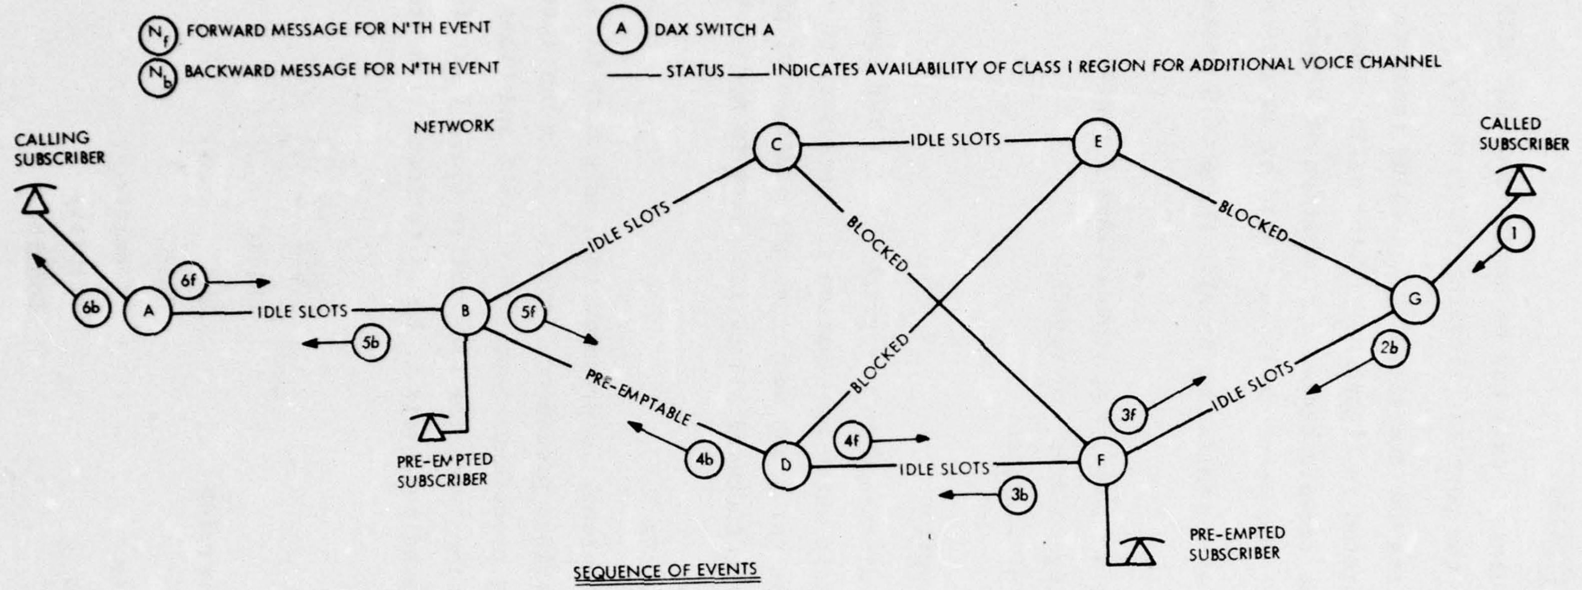

SEQUENCE OF EVENTS

1. CALLED SUBSCRIBER GOES OFF-HOOK
2. b - ALLOCATE F-G CHANNEL,
3. b - ALLOCATE D-F CHANNEL/f - F-G CHANNEL ALLOCATED
4. b - PRE-EMPT B-D CHANNEL FOR REUSE/
   f- D-F CHANNEL ALLOCATED, PRE-EMPT OTHER CHANNEL NOT FOR REUSE
5. b - ALLOCATE A-B CHANNEL/f - B-D CHANNEL PRE-EMPTED FOR REUSE
6. b - A-B CHANNEL ALLOCATED, RINGBACK REMOVED, TWO-WAY CONVERSATION ENABLED.

Figure 3-2. CCIS Procedures for Channel Allocation Phase

3.1.4.2 Breaking Down Class I Calls

An established Class I call may be broken down due either to pre-emption or to one of the parties releasing (hanging up).

Figure 3-3 depicts the overall CCIS signaling sequence used to break down the call illustrated in Figure 3-1 if the called subscriber hangs up first. Note that the forward/backward convention of Figure 3-3 is reversed from Figure 3-1 since the sequence is initiated by an action at DAX G.

As in the case of establishing a call, three CCIS messages are required per link:

a. Request to drop a channel and repack Class I region
b. Response agreeing to drop and repack
c. Acknowledgment of response.

3.1.4.3 Precedence and Pre-emption

The procedures discussed in Paragraph 3.1.4.1 for establishing a Class I call included a link in which pre-emption was necessary to route the call over that link. This paragraph addresses the question of precedence and pre-emption during the backward allocation phase in more detail.

3.1.4.3.1 Precedence Levels

The DAX network is presumed to handle both Class I and Class II traffic of various precedence levels. It is our intent to equate corresponding precedence levels of Class I and Class II traffic in establishing the rules for pre-emptive Class I traffic. Assume these corresponding levels to be as follows: (in order of descending priority)

Class I	Class II
-----	CRITIC
Flash Override	ECP
Flash	Flash
Immediate	Immediate
Priority	Priority
Routine	Routine

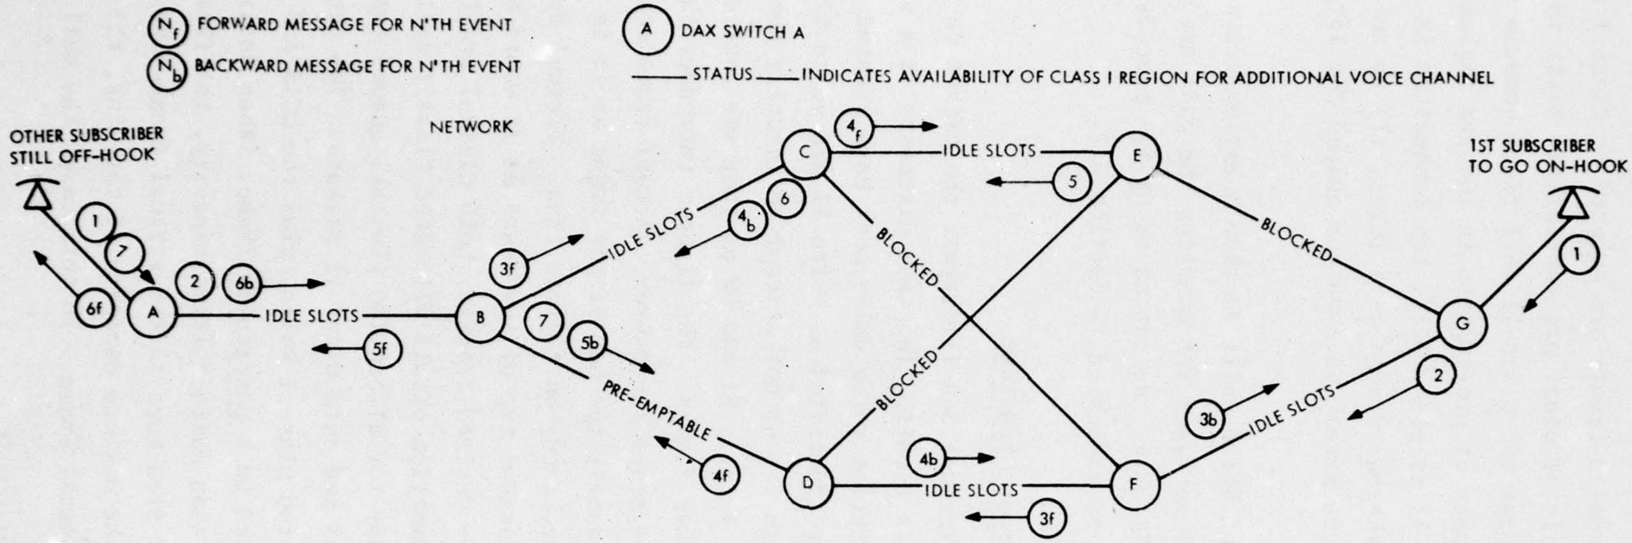

SEQUENCE OF EVENTS

FORWARD CHANNEL RESERVATION

1. SUBSCRIBER GOES ON-HOOK
2. f - DROP F-G CHANNEL
3. f - DROP D-F CHANNEL/b - F-G CHANNEL DROPPED, CLASS I REGION COMPACTED
4. f - DROP B-D CHANNEL/b - D-F CHANNEL DROPPED, CLASS I REGION COMPACTED
5. f - DROP B-D CHANNEL/b - B-D CHANNEL DROPPED, CLASS I REGION COMPACTED
6. f - BUSY TONE TO OTHER SUBSCRIBER/b - A-B CHANNEL DROPPED, CLASS I REGION COMPACTED
7. SUBSCRIBER GOES ON-HOOK, BUSY TONE REMOVED

9485-75E

Figure 3-3. CCIS Procedures for Breaking Down a Call

Although the precedence levels are equated in Class I and Class II, in cases where calls of equal pre-emptibility exist in both classes, we intend to pre-empt the Class II call first because of the difference in the impact of pre-emption in the two regions. When a voice call (Class I) is pre-empted, the connection is broken and the call is dropped, whereas a data call (Class II) is not dropped, but the data packets are merely delayed and queued for later transmission.

In the case of voice calls another question arises due to the fact that the CCIS messages for handling the call are themselves handled as Class II data. We intend to assign precedence levels to the CCIS messages as discussed in Section 3.2.

3.1.4.3.2 Pre-emption and Frame Structure

Figures 3-4 through 3-9 illustrate the effect on frame structure of a pre-emptive Class I call. The case discussed is a situation wherein two consecutive voice calls must be allocated in the Class I region. These are identified as Voice 51 and Voice 52. Note that Figures 3-4 through 3-9 do not attempt to identify the CCIS messages involved with voice calls 51 and 52 as they are considered to be transmitted in other frames. The figures introduce the concept of actual and virtual frame structures. Actual frame structure means what it implies, namely the map of the frame as it is actually transmitted (or received) on a given link. Virtual frame structure means the instantaneous map of the frame as it would have to be in order to carry the desired traffic, both Class I and Class II and represents the summation of all allocated Class I calls, plus all Class I calls ready for allocation plus all queued (and ready for transmission) CCIS and data Class II packets. The actual frame comprises the allocated Class I calls, plus the CCIS and Class II packets transmitted in a particular frame. When desired frame capacity is less than maximum frame capacity, the frame will be equal in size and structure to the virtual frame. When the desired traffic exceeds the maximum capacity of the link, the virtual frame will exceed the actual frame. In this case, the call either must be

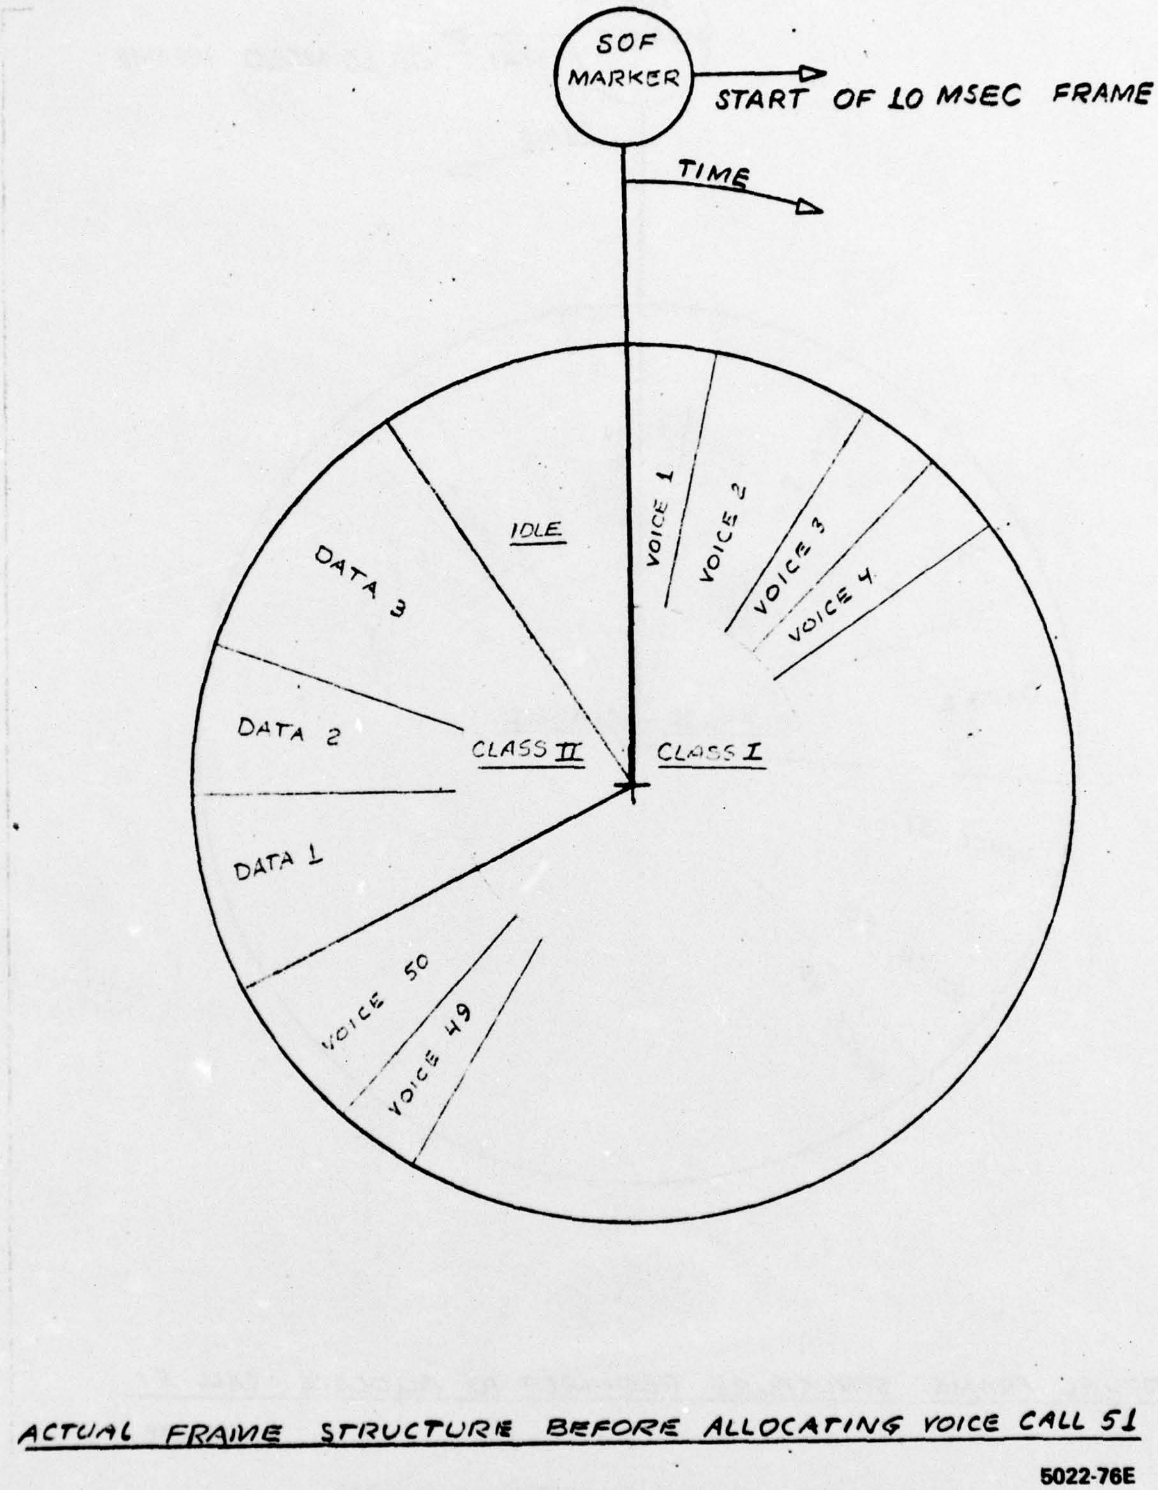

Figure 3-4. Actual Frame Before New Class I Allocation

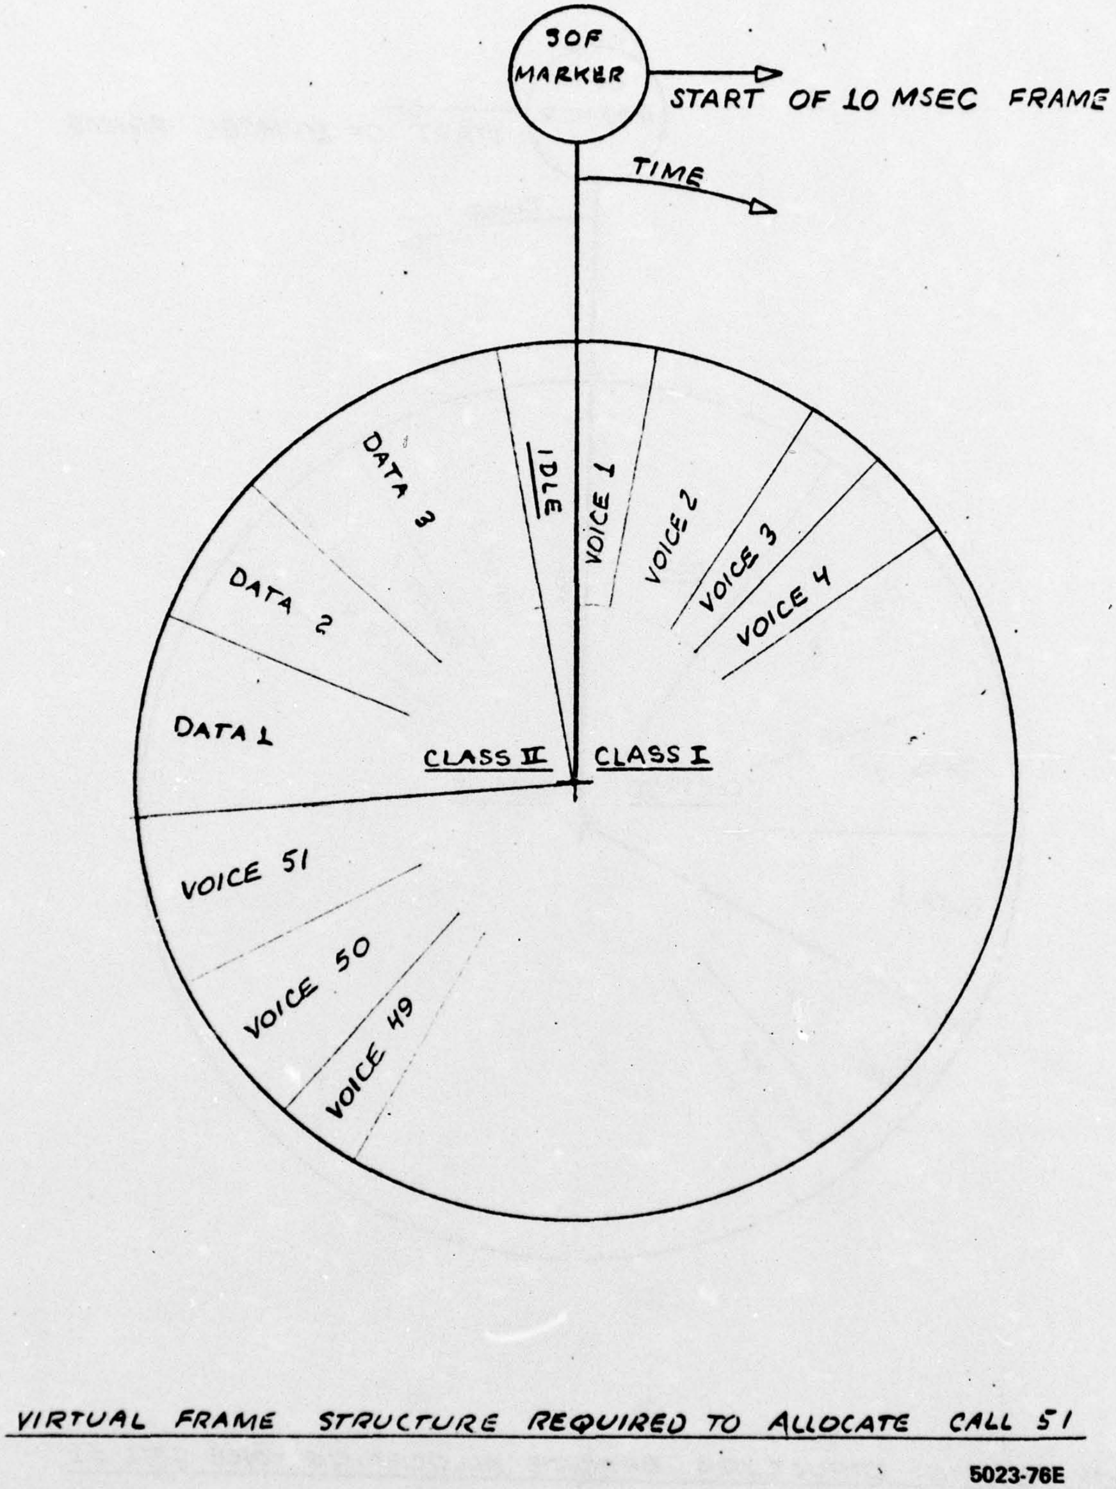

Figure 3-5. Virtual Frame Without Overlap

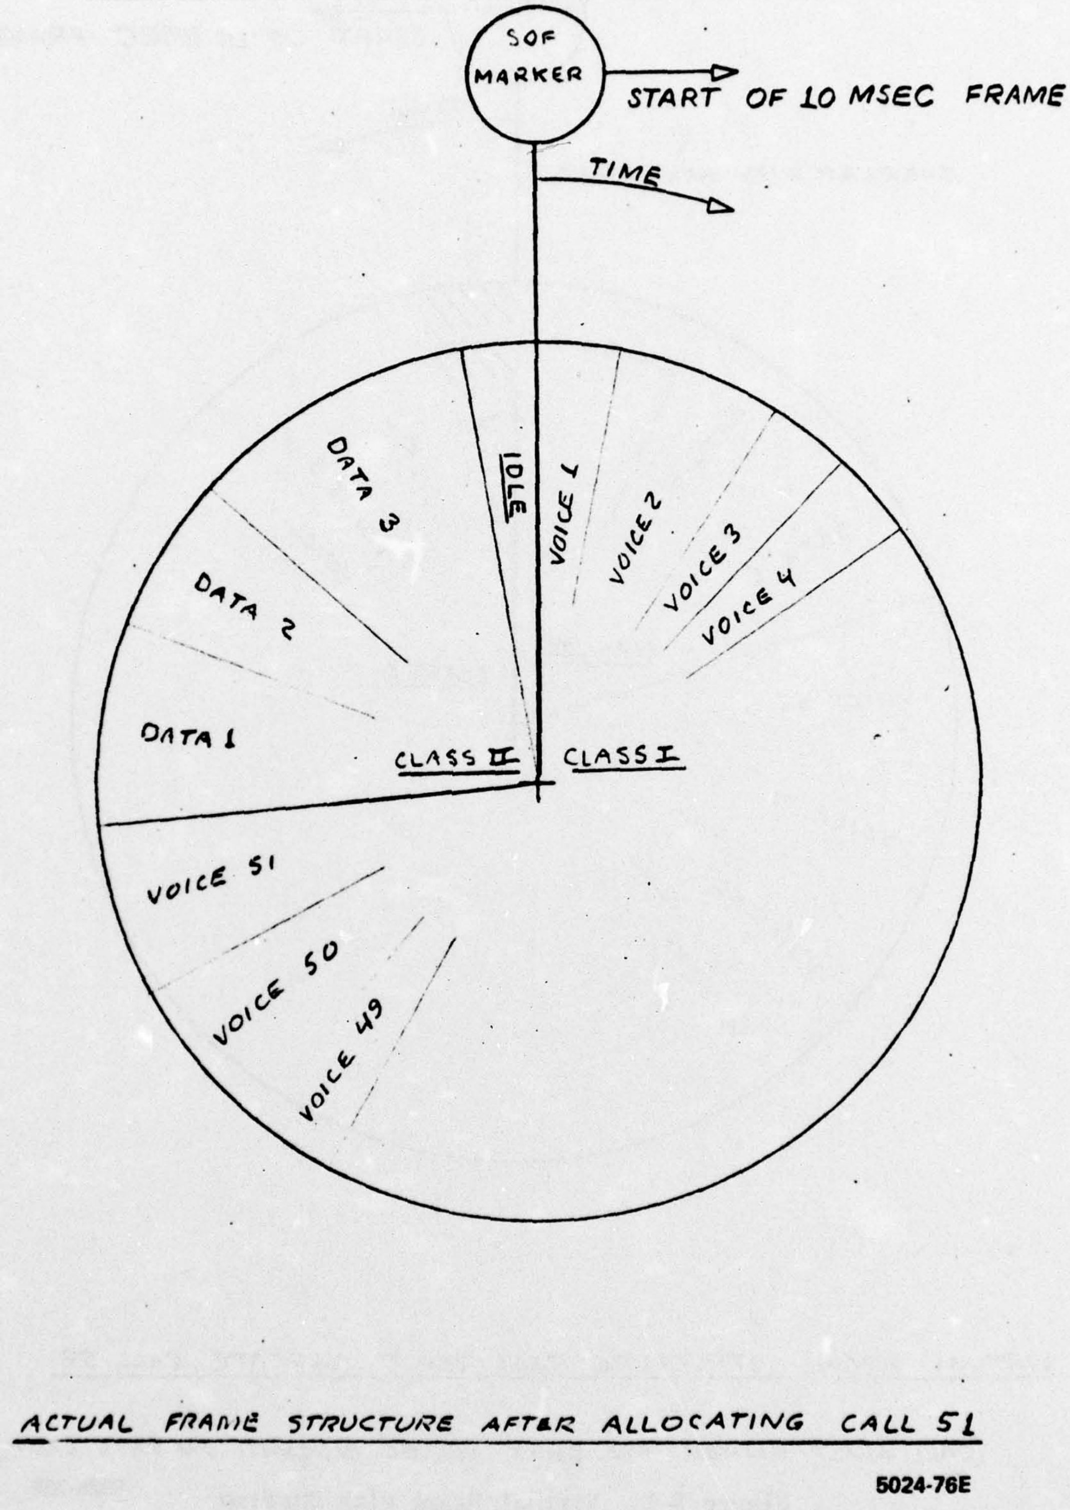

Figure 3-6. Actual Frame After New Class I Allocation

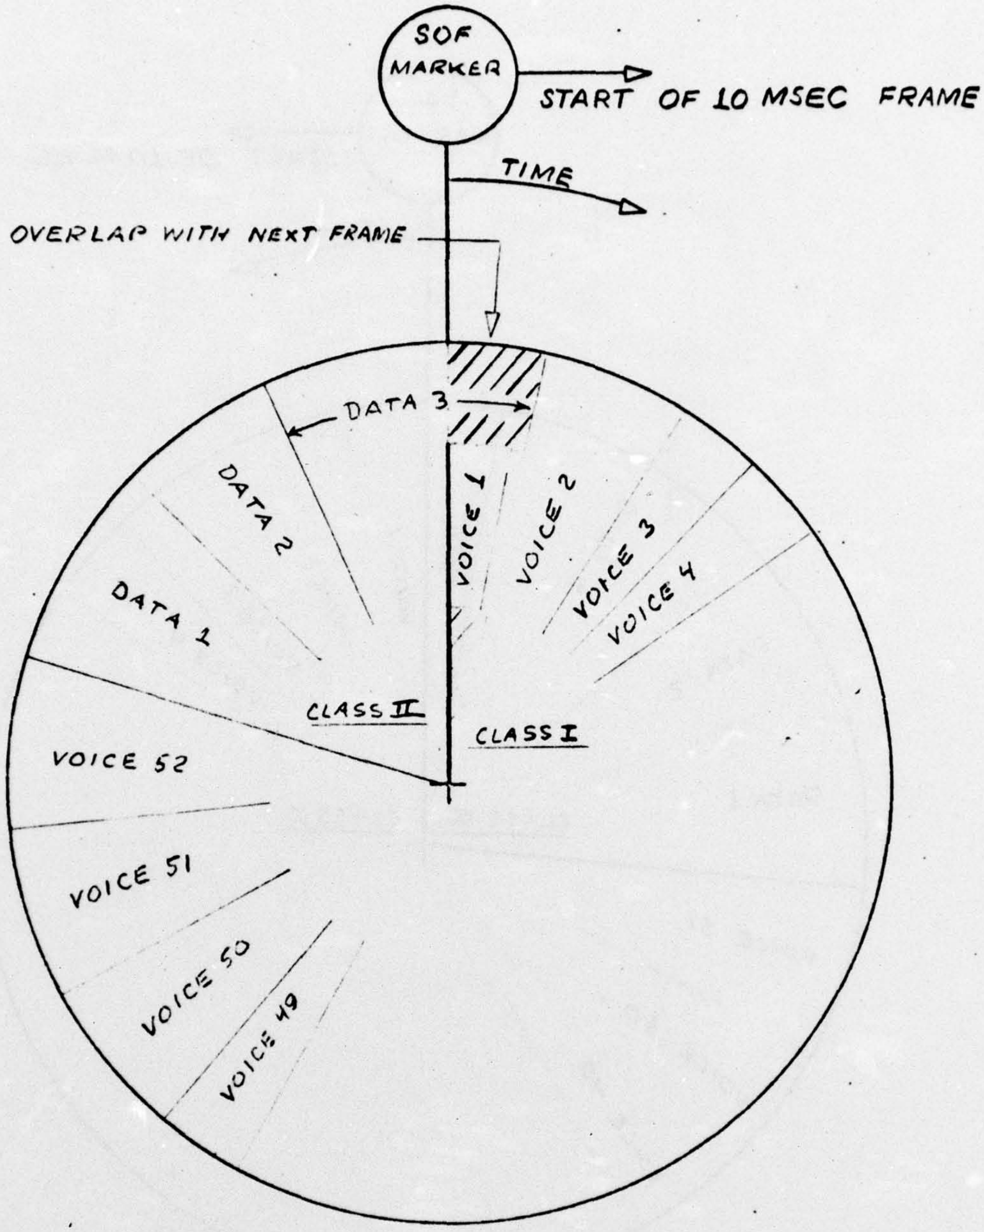

VIRTUAL FRAME STRUCTURE REQUIRED TO ALLOCATE CALL 52

(CALL MUST EITHER PRE-EMPT OR BE BLOCKED ON THIS LINK)

Figure 3-7. Virtual Frame with Overlap 5025-76E

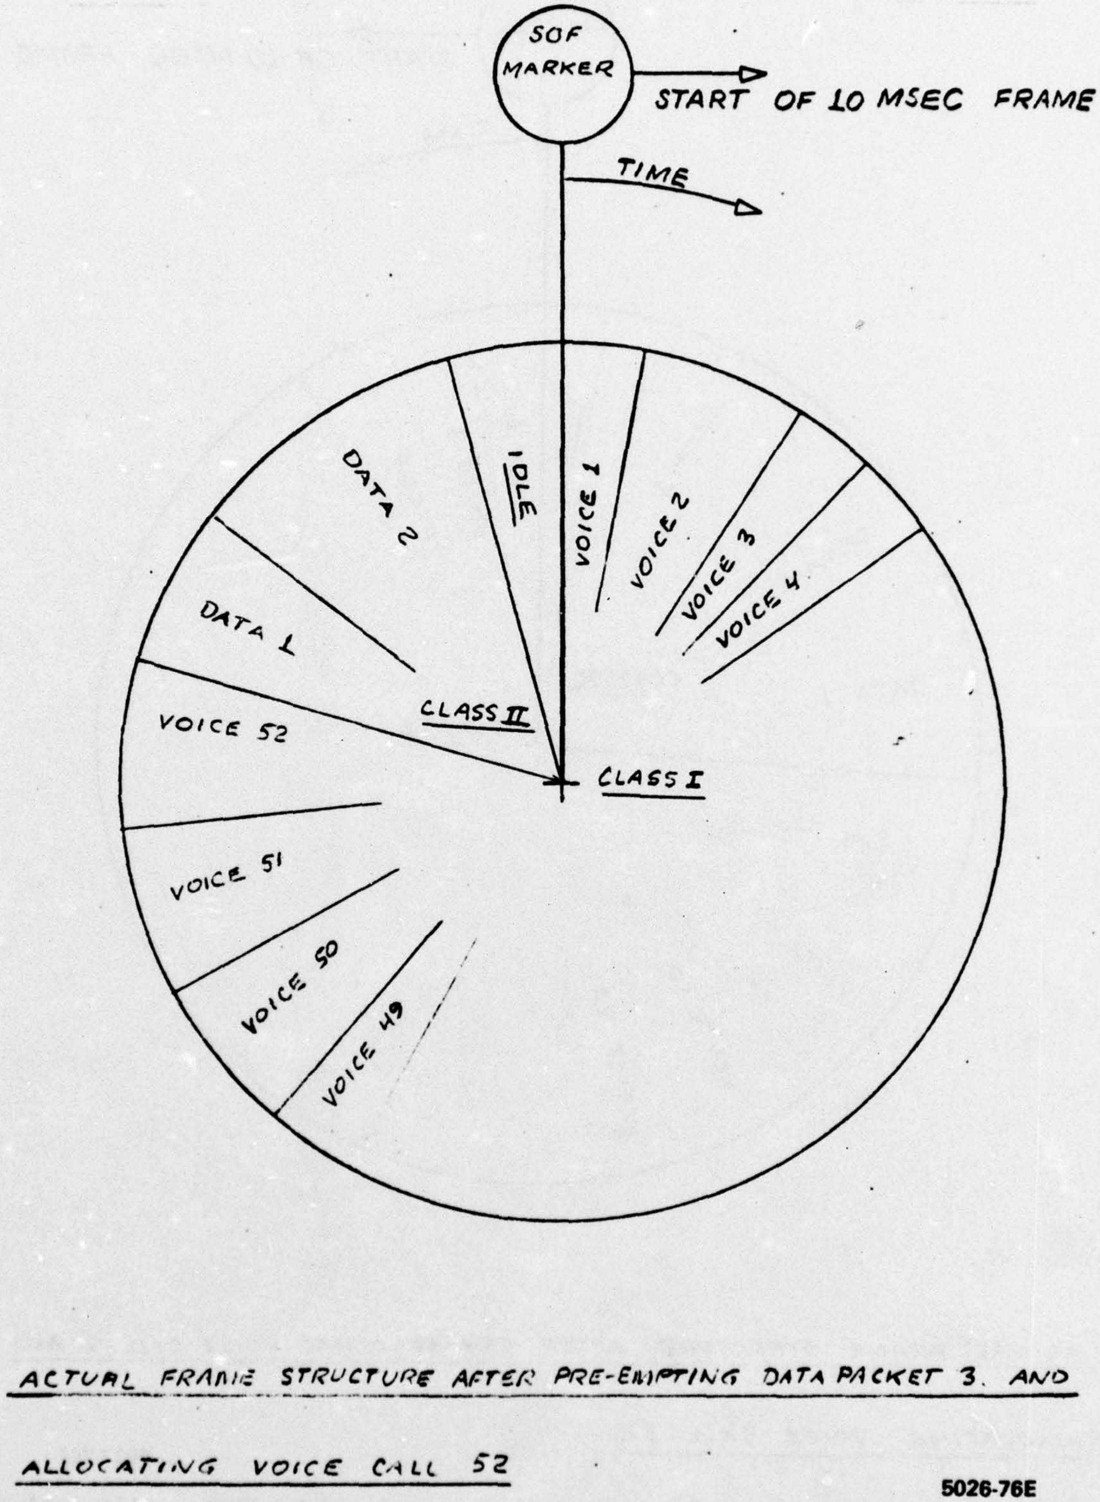

Figure 3-8. Actual Frame After Class II Pre-emption/New Class I Allocation

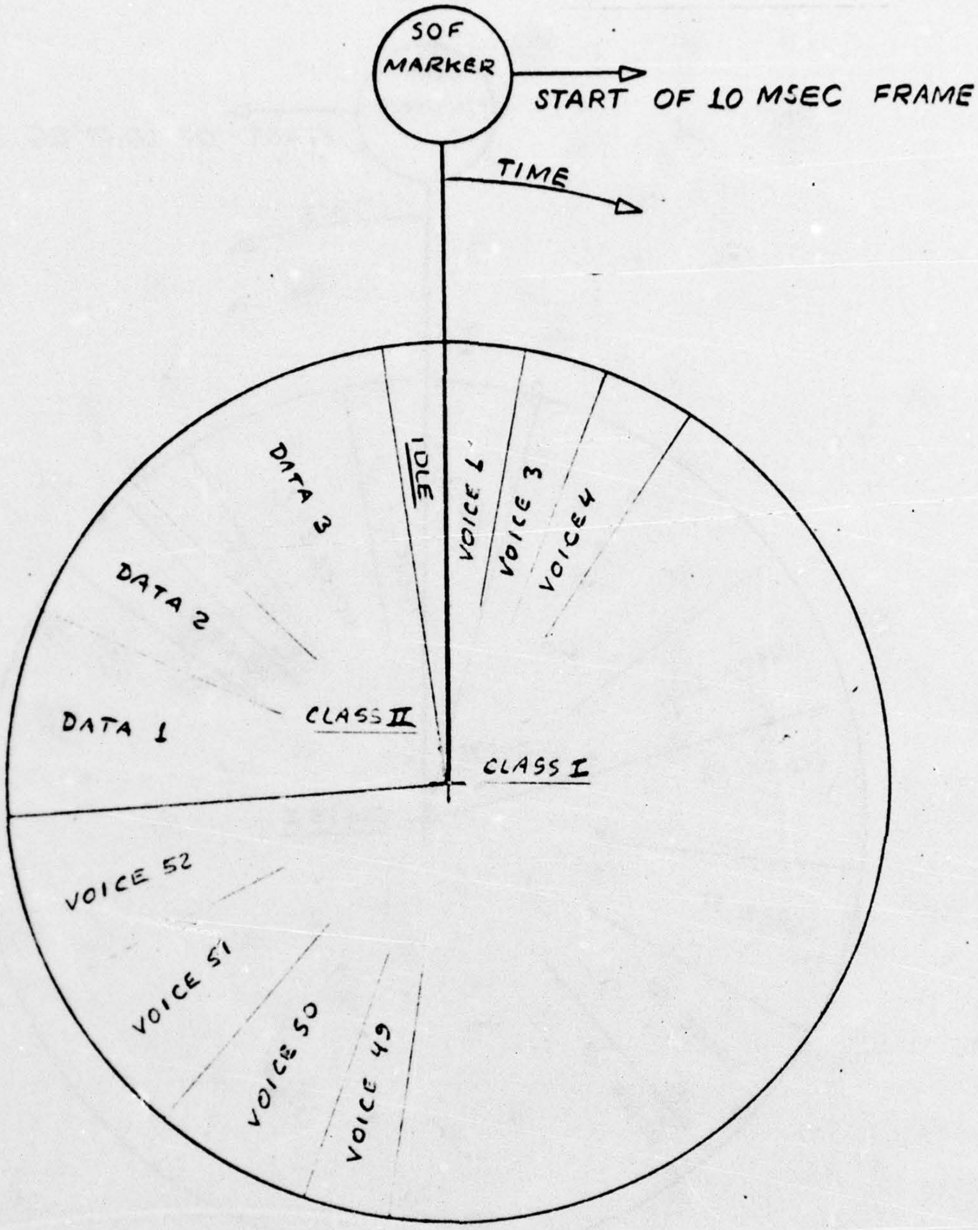

ACTUAL FRAME STRUCTURE AFTER PRE-EMPTING VOICE CALL 2 AND ALLOCATING VOICE CALL 52

5027-76E

Figure 3-9. Actual Frame After Class I Pre-emption/New Class I Allocation

blocked or must pre-empt an ongoing call (or calls) in either the Class I or Class II region.

In the example Figure 3-4 shows the actual frame structure in the frame period before voice call 51 is to be allocated, and Figure 3-5 shows the virtual frame structure required to allocate voice call 51 in the subsequent frame period. Since the virtual frame can fit satisfactorily within the 10 msec frame period the actual frame map is adjusted to conform to the virtual frame. This is shown in Figure 3-6.

Next, in a subsequent frame an allocation request is made for voice call 52, resulting in the virtual frame structure shown in Figure 3-7. It can be seen that the virtual frame exceeds 10 msec, resulting in frame overlap. Therefore, unless the new call can pre-empt an existing call either in the Class I or Class II region, it must be blocked. If the call were to be blocked, the virtual frame would revert to the configuration of Figure 3-6

Figure 3-8 depicts the virtual and actual frame structures that would result if data call 3 were to be pre-empted to make room in the frame for voice call 52. If, instead, voice call 2 were to be pre-empted, the resulting virtual and actual frame structures would be as shown in Figure 3-9.

3.1.4.3.3 Selection of Calls to be Pre-empted

A general criterion for pre-emption is that pre-emption will be done in reverse precedence order. Within this framework several trade-offs may be identified.

Given the overall system objective of maximizing frame utilization, it would seem reasonable to attempt to pre-empt a call of the lowest precedence with the same baud rate as the pre-empting call. This would minimize the pre-emption of more frame capacity than actually needed to accommodate the new call and would also minimize the number of individual calls pre-empted.

There is an undesirable aspect to this approach, however, in that pre-emptibility would be inversely proportional to commonality

of equipment. The subscriber with the most unique equipment need fear pre-emption the least even though class-marked for low precedence calling privileges. Given that the least important subscribers in a network are usually the last to receive advanced equipment, this aspect may tend to make the more important subscribers more vulnerable to pre-emption.

The approach that is currently proposed is to pre-empt at the lowest precedence level starting from the beginning of the frame, without regard to the baud rate. Given the drop/insert appraoch discussed in task B3, this will result in pre-emption of the longest-established calls first.

3.1.4.4 Preventing Precedence and Security Violations

Precedence and security classification violations will be prevented by the use of classmarks. Each subscriber in the network will be assigned both originating and terminating classmarks. Originating classmarks describe the maximum precedence level at which the subscriber may initiate a call and also the maximum security classification of the call. Terminating classmarks describe the maximum security classification of calls which the subscriber may receive. Therefore, call precedence is screened at the originating switch when the call is initiated. Once the call is established however, all terminations and interswitch channels involved in the call will be protected against pre-emption at the precedence level at which the call was placed. Call security classification is screened at both the originating switch and the terminating switch. Screening at the terminating switch will be done prior to ringing the called subscriber.

Both precedence and security classification will be carried by "traveling class-marks" contained in the CCIS messages.

Inter-DAX trunks are assumed to be classmarked for unrestricted precedence and security privileges.

3.1.4.5 Route Selection for Class I Calls

In general, Class I voice calls are initiated by a subscriber dialing a directory number into the DAX on which he is homed (originating switch). The call must then be routed to the switch on which the called subscriber is homed (terminating switch). It is assumed that a uniform numbering plan will be incorporated which permits the call to be routed deterministically, rather than by a saturation search technique. A typical example might be a numbering plan, wherein the subscriber dials from 7 to 11 digits of the form:

P-NAN-YYY-XXXX

Where P is the precedence digit (e.g. Immediate)

NAN is the area code

YYY is the station code

XXXX is the individual subscriber number

The originating switch would examine the dialed number starting with the NAN (Area Code) by searching the stored area code directory table for a corresponding entry. If a match is found with the home area code, then the station code exchange directory would be searched for an entry corresponding to the YYY. If the NAN matched a foreign area code, the call would have to be routed to an adjacent switch over a trunk (link) as indicated by additional information contained in the directory table for each foreign area code. If multiple routes to the terminating switch exist (usually expected to be the case), one of the routes out of the originating switch will be designated as the primary trunk and another as the first alternate trunk and so on. Having ascertained the potential routes that may be taken, the switch must then select one. The process of selecting the route ideally will be made powerful enough to permit flexibility in selecting the routes not only on the basis of the called party directory number but also on the basis of individual call characteristics such as precedence level, security classification, baud rate, etc. For example, in one instance it may be desirable to search for idle channel capacity in the primary trunk, then in the first alternate trunk and then go back and search for a pre-emptible channel in the primary trunk. For a different call it may be desirable to search first for idle channel capacity and then a pre-emptible channel in

the primary trunk before moving on to the secondary.

3.1.4.5.1 Routing in the Forward Direction

In the forward phase of a call a potential route is established from calling to called party and capacity in the interswitch frames along the route is reserved for that Class I call. That reserved capacity must be taken into account together with all allocated Class I channels (i.e., existing on-going voice calls) in forward routing of other subsequent voice calls. The reserved capacity may be used for Class II traffic until such time that a channel is allocated for the call when the called party goes off-hook (answers the call).

The overall reserved route at the time the call reaches the terminating represents the route over which the voice connections will probably be established.

3.1.4.5.2 Routing in the Backward Direction

It is not until the called subscriber has been rung (or goes off-hook if his equipment is capable of sending answer supervision to signals to the DAX) that a channel must be allocated in the Class I region of each trunk in the route starting at the terminating switch and working backward to the originating switch, thus establishing a two-way voice path between the calling and called subscriber. Generally speaking, this voice path will follow the route which was reserved during the forward phase of the call.

The situation can be complicated by the existence of Class II data. Suppose for example that during the backward allocation one of the switches finds that it cannot allocate the new Class I channel because of a large amount of Class II data ready for transmission over the trunk (link) and that the Class II data is of equal or higher precedence level than the Class I call in question. At this point several alternatives are available to the switch:

a. The switch may delay the channel allocation in anticipation that the Class II data will be rapidly transmitted and will recover slowly;

b. The switch may alternate route the call over a different trunk not previously involved in that call;

c. The switch may block the call and send CCIS messages both forward and backward to free the channel capacity previously reserved for that call;

d. A sequential combination of the above.

If alternative d is taken, the adjacent switch (toward the called party) could then also take alternatives b and c and this sequence of events could cascade all the way back to the terminating switch at which time the call would have to be blocked and blocking tones returned to both called and calling subscribers.

3.2 TRAFFIC CONTROL PROCEDURES

3.2.1 Problem

The Class II region of a master frame is used to carry both packetized CCIS messages and packetized data traffic such as interactive data and bulk data transfers. These packets, differ from the "voice channels" of the Class I region, in many respects. The major differences are:

a) The switched channels of the Class I regions are established on a call basis, i.e., a fixed number of bits are transmitted in the Class I region of every frame from the time the call is established to the time the call is broken down. In contrast, the data packets of the Class II region are established on an as-needed basis, i.e., as data messages are received a packet is formed and the required number of bit slots is assigned in the Class II region on a First-In/First-Out by precedence basis for only that frame in which the message is to be transmitted.

b) The Class I switched channels contain no link or network control information. The bits in each channel pass through the switch unaltered. (The control information is contained in CCIS Message packets transmitted in the Class II region.) The Class II data packets (including CCIS packets) are self contained in that each packet includes the necessary link and network control information for handling that packet.

The problem addressed in this section is to analyze the various functional requirements and define a set of control procedures which will permit efficient intermingling of data traffic having mixed precedences and security classifications without violation in either area.

3.2.2 Objectives

To define a set of data traffic control procedures which will permit:

1. Accurate transfer of packets across a link in a manner which insures detection of and recovery from transmission errors.
2. Distinguishing CCIS packets from data packets.
3. Routing of data packets towards their ultimate destination.
4. Minimizing the cross-network delay based on precedence level of each data packet so that highest precedence data has the least delay.
5. Handling of Class II traffic of mixed precedence levels.
6. Pre-emption of lower precedence calls (Class I and Class II) by pre-emptive data traffic.
7. The prevention of precedence and security classification violations.

3.2.3 Approach

The approach towards meeting the objectives stated above is discussed below.

3.2.3.1 Accurate Transfer of Packets Across a Link

In order to prevent the loss of data due to transmission errors, a positive intra-link signalling scheme is required which provides for detection of transmission errors and a means of receiving from such errors, either through forward error correction (FEC) techniques, automatic-repeat-request (ARQ) techniques or a combination of both. The ADCCP packet format which has been proposed for use in the DAX (Ref Sec. 2.1) permits an ARQ capability which will satisfy this objective.

3.2.3.2 Distinguishing CCIS Packets from Data Packets

Because CCIS packets are processed differently by a DAX than are data packets, a DAX must be able to distinguish between the two types of packets. In order to reduce the processing load, it is desireable to make the recognition process as simple as possible. This will be true regardless of whether a distributed or centralized processing architecture is implemented. The recommendation was made in Sec. 2.4 to use a bit in the packet address field as a CCIS flag, and that recommendation is also made here.

Another way in which CCIS packets might be distinguished from data packets is by their relative positions in the Class II region. In order to minimize the number of packets headers that are examined in searching for CCIS messages, it has been suggested that a superprecedence level be assigned to the CCIS packets and that they be located (continguously, in the case of multiple CCIS messages) in the beginning of the Class II region. Thus when the first non-CCIS (data) packet is encountered, it could be assumed that there were no more CCIS packets from the frame, thereby minimizing the search procedure. It is conceivable that in some cases such a scheme could result in a one frame delay in transmitting a critic data packet out of the next node due to the timing relationship of the incoming and outgoing trunks. Such a delay might or might not be tolerable.

A refinement might be to assign the precedence level of the associated Class I call to the CCIS packet when determining the availability of capacity in the frame and if capacity is available to assign an artificial precedence level (perhaps FLASH) to the packet to determine where in the Class II region the message is placed. This scheme might represent an acceptable compromise between minimizing network delays of some intermediate-level data packets and minimizing processing overhead.

In summary, CCIS packets will be made identifiable by a bit in the second field of the ADCCP format with possible use of relative location in the Class II region to identify the CCIS packets in a frame.

3.2.3.3 Routing of Data Packets

As in the case of Class I switched traffic, Class II packets must also be routed through the network from the originating (or source) switch to the terminating (or destination) switch. This does not apply, of course, to calls between two subscribers homed on the same switch.

The routing function for Class II traffic differs from that for Class I traffic in two important aspects discussed below.

3.2.3.3.1 Routing of Class I traffic is accomplished during the Call Initiation/Channel Allocation Phase, and all subsequent conversation or data travels over the same route through the network, i.e. over the same interswitch links, for the duration or the call. Routing of Class II traffic is done on a per-packet basis, i.e. each packet may be routed independently of other packets associated with the same call.

The routing scheme proposed for Class I traffic (Section 3.1) incorporates an alternate routing capability accomplished through the use of routing tables stored at each switch. These tables define primary and alternate routes from that switch to every other switch in the network. A Class I call will always be routed over the primary route if possible, then the first alternate route, etc. The routing table at each switch is unique to that switch and is fixed (i.e., the primary route to a given destination switch is always the primary route, the first alternate is always the first alternate, etc.).

The routing scheme proposed for Class II traffic will incorporate an adaptive routing capability somewhat equivalent to the alternate routing for Class I traffic. In this case however the Class II packets (except for CCIS) will be directed to the route which has the smallest predicted delay from the switch in question to the destination switch. The predicted delays to each destination switch will be contained in a table which is updated periodically. A technique being considered is the routing algorithm as it exists in the ARPA network today.[33],[46]

This technique which is still in the evolutionary stage might require further modifications or refinement to accommodate the precedence levels handled by DAX. For example, the ARPA algorithm uses queing delays as a factor in calculating predicted overall delays. In a network carrying data traffic of various precedence levels, the actual queing delay would be dependent on the procedence level of the packet being routed. It is not known at this time whether the algorithm should be expanded to include this factor or whether it could be safely ignored. Further analysis is necessary.

3.2.3.4 Minimizing the Cross-Network Delay Based on Precedence Level

Data packets will be serviced and retransmitted on a FIFO basis within precedence levels. A packet queue will be constructed for each outgoing trunk and will be emptied out (transmitted) in order of decreasing precedence. The queue will be represented by a linked list for each precedence level in which each entry will represent one data packet. The last entry of the list for each precedence will link to the first entry of the list for the next lower precedence level. CCIS packets will have precedence levels assigned as described in Paragraph 3.2.3.5.2, and thus will be merged into the Class II region of a frame after data packets of higher precedence level and before data packets of equal or lower precedence level.

3.2.3.5 Class II Precedence Levels

As mentioned earlier, the Class II region can contain two types of traffic; data packets and CCIS message packets. The question of precedence and pre-emption for each type of packet is treated separately in the following paragraphs.

3.2.3.5.1 Data Packet Precedence

Since interswitch data packets are the vehicle for transferring the actual call data as well as the interswitch/network control information, the precedence level of the data packet is always the same as the precedence level of the Class II call. It is our intent to equate corresponding precedence levels of Class I (switched) calls and Class II (Data) calls in establishing the rules for pre-emptive traffic. Assume these corresponding levels to be as follows:
(in order of descending precedence)

Class I (voice)	Class II (data)
----(none)	W (critic)
FO (Flash Override)	Y (ECP)
F (Flash)	Z (Flash)
I (Immediate)	O (Immediate)
P (Priority)	P (Priority)
R (Routine)	R (Routine)

3.2.3.5.2 CCIS Message Packet Precedence

CCIS packets are the vehicle for transferring only interswitch/network control information about a Class I call. The actual call data is carried in the Class I region. Thus, two separate precedences must be considered:

a) The precedence level at which Class I transmission capacity is reserved, allocated, and protected from being pre-empted. This is called the natural precedence level and is the level at which the subscriber initiates (dials) the call.

b) The precedence level used in transmitting to another DAX, the CCIS messages associated with the Class I calls. Figure 3-10 shows the precedence levels that will be assigned to CCIS messages as a function both of dialed precedence levels and the type of CCIS message to the transmitted. The rationale behind Figure 3-10 is discussed below.

Let us start with the assumption that in general CCIS message packets associated with a Class I call should carry the natural precedence of the call as initiated by the calling subscriber. This can be justified by reasoning that the signalling for a Class I call is neither more nor less urgent than the call itself. Now let us further assume that if a CCIS message does not initiate an action that will result in modifying the map of the Class I region that Class I calls should in general never be pre-empted in order to transmit that CCIS packet, but that Class II packets may be pre-empted if they are of lower precedence than the CCIS packet. This is reflected in Figure 3-10, columns 3 and 6 of which show that the natural precedence levels assigned to those CCIS messages apply to pre-empting Class II packets only; the CCIS messages are of routine precedence with respect to the Class I region.

Dialed Class I Precedence	Equivalent Class II Precedence	Call Initiate Msgs. (Reserve Capacity) Note 1	Answer Back Msgs. (Allocate Channel) Note 2	Call Release Msgs. (Drop Channel) Note 1	Other Msgs. (No Effect On Frame) Note 1
--	W	-	-	-	-
FO	Y	Y	Y	Z	Y
F	Z	Z	Z	Z	Z
I	O	O	O	Z	O
P	P	P	P	Z	P
R	R	R	R	Z	R

Note: 1 Indicated precedence level applied to class II region only. (Class I calls will not be pre-empted to transmit CCIS message)

2 Indicated precedence level applied to Class I and Class II regions. (Either Class I or Class II calls may be pre-empted to transmit CCIS message)

Figure 3-10. Precedence Levels Assigned to CCIS Messages

Next, consider the type of CCIS message that initiates an action that will result in allocating a new channel in the Class I region, namely answer back messages. It is proposed to assign the natural precedence level of the call to the answer back messages, pre-empting in the Class I region if necessary. Since the weighted Class I channel size is on the order of 150 bits and the answer-back CCIS message packets is on the order of 120 bits, in most cases the Class I call that is pre-empted for transmitting the CCIS message would have to be pre-empted to fit the new channel into the Class I region anyway.

Finally, consider call release messages, which result in Class I channels being dropped. In this case, the natural precedence level of the released call is meaningless because the call has terminated anyway. It is proposed to assign an arbitrary level of FLASH to call release messages, to apply to the Class II region only, so as to facilitate DAX retrieval of transmission capacity as quickly as practical.

A unique situation can arise when the Class I traffic fills the entire frame to the point where the Class II region is not large enough to contain a CCIS message. In such a case the system would be locked up forever unless some provision were made to transmit a CCIS message to handle release calls. This situation is called the lock-up condition and is discussed in Paragraph 3.6.

3.2.3.6 Pre-emption

This paragraph addresses the problem of handling pre-emptive Class II traffic. Pre-emptive data packets are discussed in 3.6.1; pre-emptive CCIS message packets are discussed in 3.2.3.6.3.

3.2.3.6.1 Pre-emptive Data Packets

A data packet of a given precedence level may pre-empt packets of lower precedence levels or voice calls of lower precedence levels (Ref. 3.2.3.5.1). However, pre-emption of a lower level data packet may cause that data packet to initiate pre-emption of a voice call in a subsequent frame. This situation is illustrated in Figure 3-11 which depicts a frame in which the Class II region is not large enough to include all data packets which are queued for transmission. The data packet queue to the right hand side of the figure contains the linked lists representing the Class II data packets available for transmission at the first opportunity. Within

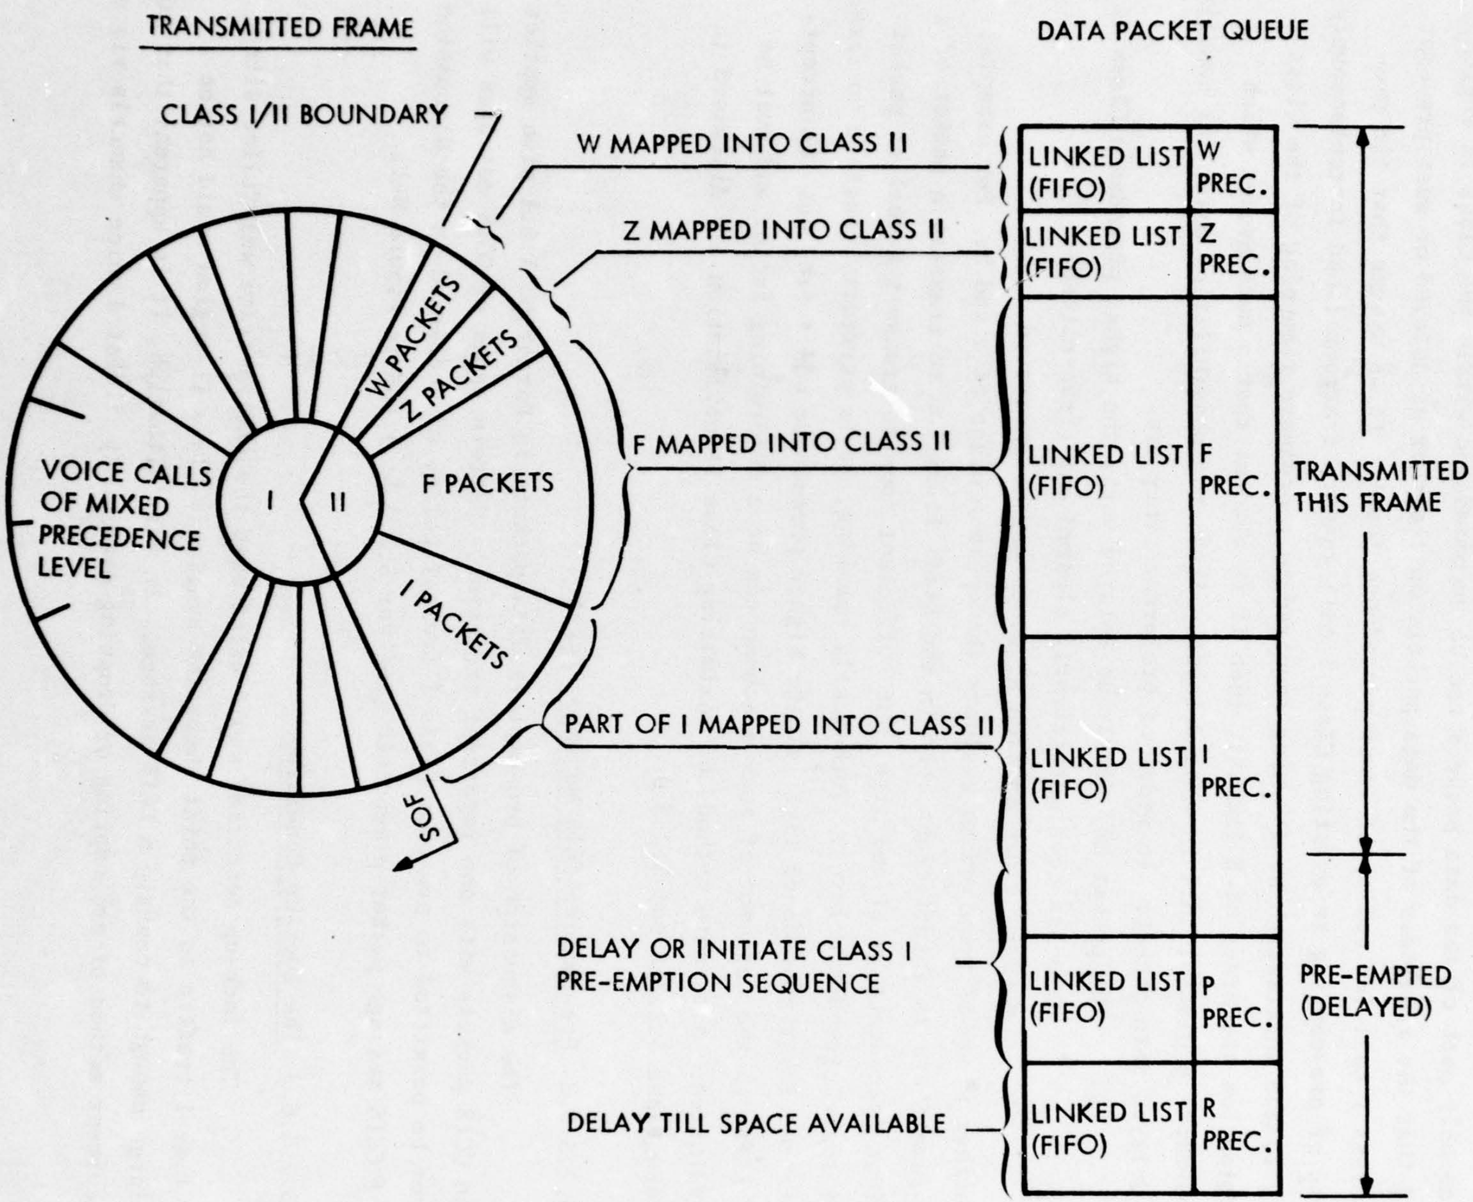

Figure 3-11. Handling of Pre-Emptive Class II Data Packets

each linked list, the data packets are time-ordered on a FIFO basis. Within the queue, the linked lists are ordered by precedence from highest to lowest.

The transmitted frame is shown to the left of Figure 3-11. It can be seen that only part of the data packets can be accommodated within the Class II region, and that the remainder of the data packets must either be delayed or must pre-empt a Class I call (or calls) of lower precedence level. If we assume that the process of pre-empting an existing Class I call cannot be accomplished instantaneously due to the interswitch handshaking required to coordinate remapping of the Class I region (on the order of 8 frames), then it is evident that a data packet which pre-empts a Class I call is also delayed. Thus, two possibilities exist for handling a delayed data packet. In order of preference they are:

a. The packet may merely be delayed until the higher precedence Class II data has been sufficiently cleared out (Transmitted).

b. A Class I pre-emption sequence may be initiated.

Further study is required to evaluate the optimum usage of a and b. For example, assume there is insufficient room in the Class II region to transmit a packet of a given precedence level and size, but sufficient room to transmit a smaller packet of lower precedence level. Technically speaking, it is perfectly feasible to transmit the lower precedence first if the higher precedence won't fit, but the acceptability to the customer of such a scheme can be a determining factor and must be evaluated. Alternate methods of maximizing Class II utilization are discussed in Paragraphs 3.2.3.8 and 3.2.3.9.

3.2.3.6.2 Pre-emptive CCIS Message Packets

The discussion of pre-emptive data packets in Paragraph 3.6.1 also applies to CCIS packets with one important exception. Certain types of CCIS messages will not be permitted to pre-empt Class I (voice) calls, as indicated in the discussion of CCIS message packet precednece (Ref Par 3.2.3.5.2) and in Figure 3-11.

3.2.3.6.3 The Lock-Up Condition

The lock-up condition would exist when the entire frame was filled with Class I traffic to the point where the available Class II region would not be large enough to contain a CCIS message. In this situation, it is apparent that our current method of pre-empting or dropping previously allocated voice channels via an

exchange of CCIS messages across the trunk in question will not always work. While it is true that quasi-associated CCIS messages could be exchanged over alternate routes, there is not quarantee that an alternate route will always be available. Therefore, either some alternate method of transmitting CCIS messages must be developed, or else the lock-up condition must be avoided in the first place. The method of avoiding the lock-up condition is simple enough--never transmit a CCIS message that will result in the entire frame being allocated to Class I traffic.

An alternate method of transmitting CCIS messages during a stress condition has been considered. The CCIS message could be transmitted in the Class I region by momentarily interrupting an appropriate number of contiguous voice channels, preferably of lower precedence levels than the CCIS messages. The affected voice calls would appear to the associated subscribers as noisy channels for one frame interval (10 msec) but the calls would not be dropped. This method conceivably could be extended to data packets as well, thus making it unnecessary for a Class II data call to ever permanently pre-empt (drop) a Class I voice call.

Transmitting Class II packets in the Class I region is a possibility but incurs certain drawbacks. The drawbacks relate to the problem of recognizing a data packet contained in the Class I region by detecting the flag sequence and the problem of rejecting false alarms caused by simulation of the flag sequence by the random bit stream of the Class I voice calls. The first problem results in added processing complexities at the receiving DAX (software and possibly hardware as well). The second problem results in added processing load (time burden) at the receiving DAX. For these reasons, this approach is not recommended at this time, although it can be given further consideration in the future, if this seems desirable.

3.2.3.7 Establishing a Class II Call

This paragraph describes a procedure for establishing a typical terminal to computer Class II call through a hypothetical network. The network connectivity and link availability was chosen to be the same as the example used in Figure 1 of Section 3.1 to facilitate a direct comparison of Class I and Class II procedures. The sequence of events is shown on Figure 3-12 but a few words here are in order.

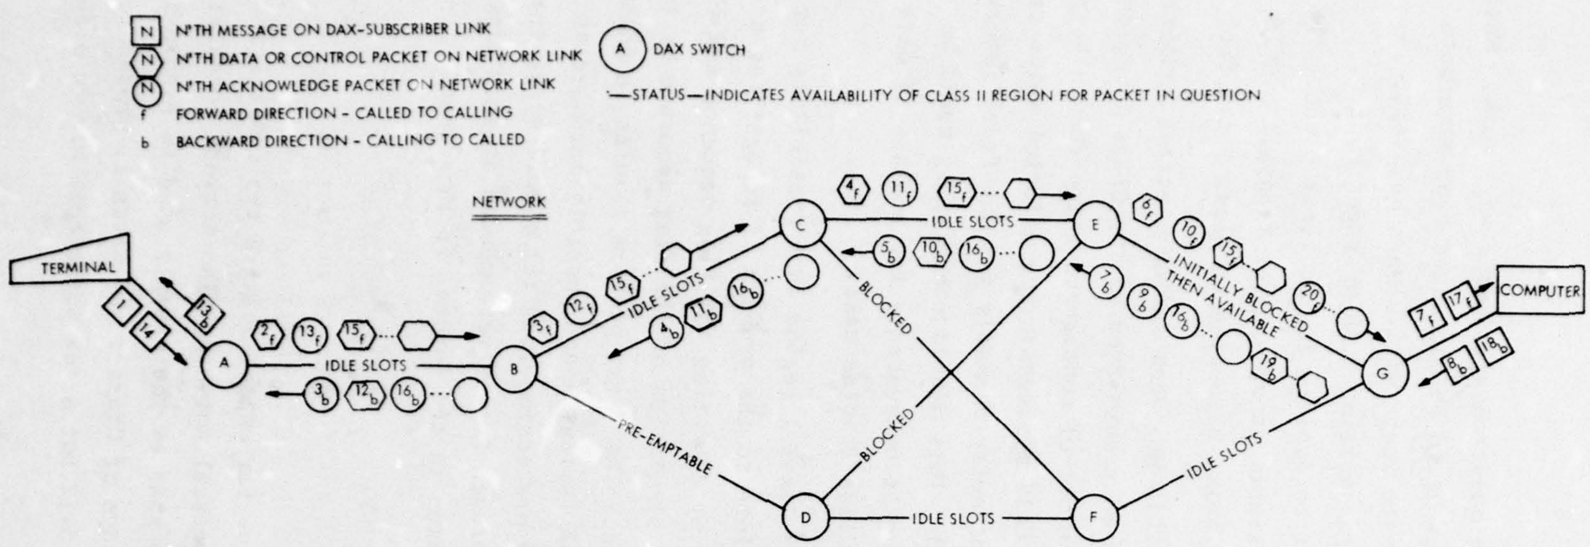

SEQUENCE OF EVENTS

1	1	TERMINAL INITIATES CALL TO COMPUTER/DAX A ALLOCATES BUFFER TO TERMINAL, TRANSLATES DIALED NO., FINDS COMPUTER IS HOMED ON DAX G
2	2_f / 3_b	DAX A INITIATES CALL INITIATE PACKET TO DAX B/DAX B ACKNOWLEDGES RECEIPT OF CALL INITIATE PACKET
3	3_f / 4_b	DAX B ROUTES CALL INITIATE PACKET TO DAX C/DAX C ACK'S PACKET
4	4_f / 5_b	DAX C ROUTES CALL INITIATE PACKET TO DAX E/DAX E ACK'S PACKET
5	6_f / 7_b	DAX E QUEUES PACKET UNTIL E-G LINK CAPACITY AVAILABLE, ROUTES PACKET TO DAX G/DAX G ACK'S PACKET
6	7_f / 8_b	DAX G SENDS CALL INITIATE MSG TO COMPUTER, ALLOCATES BUFFERS TO TERMINAL AND COMPUTER/COMPUTER ANSWERS CALL
7	9_b / 10_f	DAX G INITIATES ANSWER PACKET TO DAX E/DAX E ACKNOWLEDGES RECEIPT OF ANSWER PACKET
8	10_b / 11_f	DAX E ROUTES ANSWER PACKET TO DAX C/DAX C ACK'S PACKET
9	11_b / 12_f	DAX C ROUTES ANSWER PACKET TO DAX B/DAX B ACK'S PACKET
10	12_b / 13_f	DAX B ROUTES ANSWER PACKET TO DAX A/DAX A ACK'S PACKET, ALLOCATES BUFFER TO COMPUTER AND TERMINAL
11	13_b / 14_f	DAX A REMOVES CALL ORIGINATING PACKET FROM ACK. QUEUE, SENDS ANS. MSG TO TERM./TERM. SENDS 1'ST MESSAGE OF CALL
12	15_f .... / 16_b	DAX A PACKETIZES MESSAGE, ROUTES PACKETS THRU NETWORK TO DAX G/EACH PACKET ACK'ED ON EACH LINK
13	17_f / 18_b	DAX G REASSEMBLES COMPLETED MESSAGE, SENDS IT TO COMPUTER, SENDS RFNM TO DAX A/COMPUTER SENDS RESPONSE MSG TO DAX G
14	19_b .... / 20_f	DAX G PACKETIZES RESPONSE, ROUTES PACKETS THRU NETWORK TO DAX A/EACH PACKET ACK'ED ON EACH LINK
	ETC.	

9366-75E

Figure 3-12. Procedures for Establishing a Class II Call

Notice that Figure 3-12 shows an acknowledge packet for every control and data packet. This is because in our example we are considering the packets only for the particular call in question. ADCCP procedures provide for summary acknowledgements of groups of packets, which would be the case if multiple calls were being handled simultaneously, but not in this example. In addition, note that the DAX E to DAX G link is blocked initially, and then subsequently becomes available. In our example, the DAX E to DAX D link is also blocked, and so DAX E delays (queues) the call initiate packet for later transmission. The call illustrated is from a terminal homed on DAX A to a computer homed on DAX G. The sequence of events described in Figure 3-12 is:

1) The calling terminal homed on DAX A initiates a call to the computer which is homed on DAX G. DAX A allocates a message buffer to the terminal, translates the dialed number, and determines that the computer is located at DAX G. DAX A then looks in its routing table and determines that the primary (minimum delay) route is over the DAX A--DAX B link. (Alternate routes out of DAX A are not shown in Figure 3-12).
2) DAX A then generates a call initiate packet addressed to DAX G, and transmits the packet to DAX B which acknowledges receipts of the call initiate packet by transmitting an acknowledge packet backwards to DAX A. The purpose of the call initiate packet is twofold:
   a) It requests the terminating switch (DAX G) to allocate buffer capacity to the call and to inform DAX A when the buffer has been allocated.
   b) It permits the originating switch (DAX A) to determine the availability of the called computer before accepting a message from the terminal. (The computer may be down or otherwise unavailable.)
3) DAX B then examines the address contained in the call originating packet, indexes into its routing tables, finds the primary route is to DAX C, and transmits the packet to DAX C which subsequently acknowledges receipt of the packet. DAX B then discards the packet.
4) DAX C in turn transmits the packet over the primary route to DAX E and, after receiving an acknowledgement from DAX E, discards the packet. Note that the tandem DAXs only need retain the packet long enough to determine that the packet has been received correctly by the next switch. Message accountability is maintained only by the originating and terminating DAXs (A and G).

5) DAX E finds the primary (E-G) link blocked, and queues the call originating packet for subsequent transmission. Eventually, frame capacity becomes available for the queued packet and it is transmitted to DAX G which acknowledges receipt of the packet. DAX E then discards the packet.

6) DAX G detects that the packet is addressed to itself and finds the dialed computer number in its local directory. DAX G then signals to the computer and receives an answer. Upon determining that the computer is available, DAX G allocates the buffers as required by the call initiate packet.

7) DAX G then generates an answer packet addressed to DAX A confirming both the buffer allocation and the availability of the computer for receiving a call. The answer packet is then transmitted over the primary route, in this example to DAX E.

8) DAX E routes the answer packet to DAX C, receives an acknowledgement, and discards the packet.

9) DAX C routes the answer packet to DAX B, receives an acknowledgement, and discards the packet.

10) DAX B routes the answer packet to DAX A which acknowledges receipt of packet to DAX B and notes the buffer allocation to DAX G.

11) DAX A then removes the original call originating packet from its acknowledge time-out queue and sends an answer message to the terminal, authorizing the terminal to send the first message.

12) Upon receipt of the first message from the terminal, DAX A reforms the message into the required number of packets (≥ 1) and routes the packets into the network via DAX B. Each packet is routed individually, and may arrive at the terminating switch (DAX G) non-sequentially and over different routes depending on traffic loading and dynamically updated routing tables throughout the network. The individual packets are acknowledged to the transmitting (tandem) DAXs as they are routed over the various links.

13) When DAX G has received all packets associated with the message, it reassembles and if necessary depacketizes the message and transmits it to the computer. After transmission to the computer has started, DAX G also transmits a request for next message (RFNM) back towards DAX A. If another message is waiting for transmission to DAX A the RFNM may be piggybacked on it.

14) When the computer response to the first message is received at DAX G, that message is packetized and sent through the network to DAX A and the call continues a message at a time as described above.

3.2.3.8 Partial Packets

Our approach thus far to transmission of Class II traffic has been to transmit only whole packets in any given frame and to pad the remainder of the Class II region with idle bit sequences. This approach has no adverse effect on frame utilization as long as all of the Class II data and CCIS packets ready for transmission will fit into the current frame. In cases where there is not sufficient capacity in the Class II region of a given frame to hold all of the queued packets, some frame capacity will be wasted because only rarely will the number of available bits be exactly equal to an integral number of packets. These wasted bits therefore will decrease average frame utilization under heavy traffic conditions. Even worse is the fact that under some conditions of heavy Class I traffic, some long packets may experience inordinately long queuing delays due to long holding times of Class I calls.

A concept has been formulated which would permit transmission of partial packets. Under this concept, any residual Class II capacity too small to contain the next whole packet would be used to transmit a portion of the next packet. At a minimum, the partial packet would contain the flag (F), the address (A), and control (C) fields in order to simplify processing at the receiving DAX. At a maximum, the partial packet would contain F, A, C and part or all of the I fields; the frame check sequence (FCS) would not be split across two successive frames.

An example of the partial packet concept is shown in Figure 3-13. In the example shown, data (or CCIS) packets 1, 2 and 3 are queued for transmission at the start of frame N in which there is sufficient capacity only for packet 1 with some number of bits to spare. Packet 1 is transmitted in its entirety in frame N,

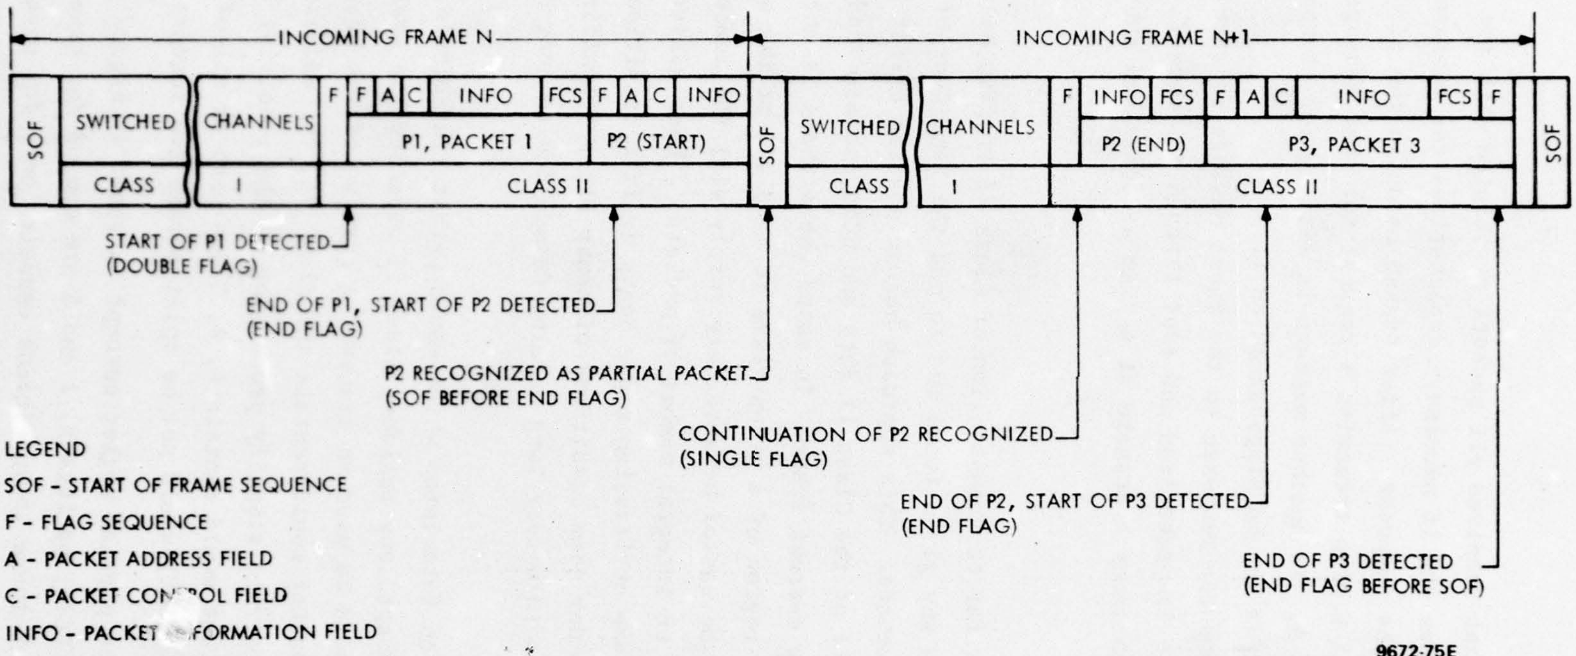

LEGEND

SOF - START OF FRAME SEQUENCE

F - FLAG SEQUENCE

A - PACKET ADDRESS FIELD

C - PACKET CONTROL FIELD

INFO - PACKET INFORMATION FIELD

FCS - PACKET FRAME CHECK SEQUENCE

Figure 3-13. Partial Packet Concept

followed by the F, A and C fields and part of the I field of packet 2, followed by the start of frame sequence (SOF) for frame N + 1. The receiving DAX will recognize that P2 is a partial packet by virtue of the fact that an SOF was detected without first seeing an end flag sequence (FCS, followed by F) for packet 2. The receiving DAX will defer processing of packet 2 and will anticipate receiving the continuation of the packet at the start of the Class II region in the N + 1 frame. When the last portion of packet 2 has been received in the N + 1 frame, the FCS is used to check the entire packet, and only then is packet 2 ACKed or ARQed, thus supporting standard link error control procedures.

As indicated by Figure 3-13 the proposed concept involves the use of a flag sequence as a delimiter between the Class I and Class II regions. This delimiter is in addition to the flag sequences used to delimit the individual packets. The reason for this is to facilitate handling of pre-emptive Class II traffic and to permit pre-emption of the partially transmitted packet 2 by a newly arrived packet of higher precedence level. Thus, if a double flag sequence were to be detected at the start of the Class II region in the N + 1 frame, the receiving DAX would discard the stored partial packet, and the transmitting DAX would have to retransmit packet 2 at the first opportunity.

It is interesting to note that in the case of extremely heavy Class I traffic the transmission of a complete packet need not be limited to two frames, but could be stretched out over many frames. Further investigation would be required to make optimum use of this characteristic.

3.2.3.9 Split Packets

An alternate approach to the partial packet concept is the split packet concept. In this approach a packet could be split into two or more smaller packets to obtain a better packing of the Class II region. Each of the smaller packets would have its own F, A, C and FCS fields. This approach might be advantageous under conditions of high transmission errors or heavy pre-emptive traffic because of reduced retransmissions, but does require more overhead, in addition to increased processor complexity. This approach is not recommended at this time but may be reconsidered in the future.

3.3 DROP/INSERT CAPABILITIES

3.3.1 Problem

The traffic-carrying portion of a master frame will be subdivided into two regions--the Class I region which carries circuit-switched data and the Class II region which carries packetized data. This task addresses the problem of controlling the Class I/Class II boundary as circuit-switched calls are established and cleared. Stated more specifically, the problem is to develop a concept for dropping and inserting circuit-switched channels in the Class I region.

3.3.2 Objectives

Any procedure for accomplishing Drop/Insert Capability must comply with the following objectives:

a. Operate with negligible error. As long as the master frame remains synchronized, the location and width of each circuit-switched channel must be known in order to prevent spurious cross-connections between subscriber calls.

b. Operate without disrupting the time transparency of the DAX to ongoing circuit-switch calls. Existing calls should not be interrupted while adding or dropping a channel or compacting the Class I region. To do so would introduce noise and risk losing crypto synch.

c. Operate with the minimum amount of transmission overhead. The average number of bits per master frame required to identify each circuit-switched channel should not be excessive.

d. Operate with the minimum amount of processor overhead. The processing resources required to maintain the Class I region map should not be excessive.

e. Operate with maximum expediency: circuit-switched calls should be inserted as quickly as practical so as to minimize connection time; on termination, calls should be dropped as quickly as practical so as to permit the system to reclaim and reassign the previously-required bandwidth. (Information capacity)

3.3.3 Approach/Progress

The current concept for providing drop and insert capabilities is based on the following assumptions which are consistent with the approaches of Section 2.1, 2.3, 2.4 and 3.1:

a. Circuit-switched channels will be continguously located in the Class I region of a master frame.
b. The Class I region will contain only circuit-switched channels, i.e. No data packets.
c. Dropping and insertion of Class I channels will be accomplished through the use of CCIS message packets transmitted in the Class II region.
d. Strict inter-DAX protocol will be observed in maintaining the Class I maps for a given link in each transmission direction. A DAX will never modify the Class I region without first notifying the adjacent DAX and receiving confirmation from the adjacent DAX of its agreement to the change.

The proposed Drop/Insert procedures are discussed below:

3.3.3.1 Channel Insertion

When a switched-network channel must be allocated to a new call, it is proposed to always add that channel to the end of the Class I region, thus displacing the Class I/ Class II boundary toward the end of the frame by a number

of bit slots equal to the number of bits (per frame) of the new channel. This approach minimizes the amount of control information that must be transferred across the link, and also minimizes the amount of internal bookkeeping at each DAX, since the Class I map need not be modified, but only extended. Note, however, that the Class I region will never be permitted to extend beyond the point at which frame lock-up would occur (Ref. Section 3.2).

In addition to specifying the width of the new channel in the CCIS Allocation Message, we propose to also include the starting address (bit position in the master frame). Although the starting address is redundant, it will be used for on-line detection of Class I map errors. The CCIS message sequence for channel insertion (allocation) is described in Section 3.1.

3.3.3.2 Channel Dropping

When a switched-network channel must be dropped, the Class I region will be recompacted at the same time the channel is dropped. The recompacting will be accomplished by subtracting the width of the dropped channel from the starting bit position of each channel in the Class I region which is later than the dropped channel. Thus the Class I region will always be compacted toward the start of frame. When two contiguous channels must be dropped, channels later than the second dropped channel will be advanced by the combined width of both dropped channels, and so on. In the case of channel dropping, only the inter-DAX channel number need be specified, but it is proposed to also include the starting address and channel width for use in on-line detection of map errors.

The CCIS message sequence for channel dropping is described in Section 3.1. It should be noted that reallocation of on-going channels due to the dropping of a terminated channel requires CCIS messages only with respect to the

terminated channel; the reallocations are accomplished automatically by both originating and terminating DAX's without necessity of additional CCIS messages.

3.3.3.3 Channel Pre-emption

In handling pre-emptive traffic it will sometimes be necessary to pre-empt (drop) one or more channels allocated to existing calls in order to make room for a channel allocated to a call of higher precedence. The procedure for pre-empting channels for reuse will be to drop the pre-empted channels and compact the frame per paragraph 3.3.3.2 and to insert the new channels at the end of the newly compacted Class I region as described in paragraph 3.3.3.1. However, a unique CCIS message will be defined for "Pre-empt For Reuse" to avoid having to exchange sequential message sequences for both the drop and then the insert functions.

For reasons of processing economy it is planned to search for channel capacity which can be pre-empted within the lowest precedence levels starting from the beginning of the frame. Given the insert and drop procedures of the preceeding paragraphs, this means that within any precedence level on a given trunk, the longest established calls will be preempted first. Other approaches are possible however, and can be examined if desired. For example, the oldest call of low-enough precedence to be preempted might not provide enough capacity for the pre-empting call, whereas the next-oldest might; the preempting algorithm might consider adequacy as well as age, within each precedence level.

3.3.3.4 Background Testing of the Class I Map

In addition to the on-line map error detection which will be done on a per-allocation and per-drop basis, it seems advantageous to consider a mechanism for exchanging Class I map information between DAX's. Such an exchange

would be routinely done in the background on a non-interfering basis. The required CCIS messages could be assigned a sub-precedence level lower than routine, meaning that they would be transmitted on a periodic basis when frame capacity is available with no other packets queued for transmission.

Since some sort of frame map update messages will undoubtedly be required for error recovery (reference Section 3.5), the background testing function could use essentially the same software mechanism.

3.4 ACCOUNTABILITY

3.4.1 Problem

The transmission concept being studied is an integrated switched telecommunications network which is to handle simultaneously, circuit switched, packet switched, and some message switched (store & forward) traffic. Accountability requirements must be identified and analyzed for each type of traffic. The problem addressed in this section is to evaluate the impact of accountability on each class of service of enroute communications from the viewpoint of:

a) Dynamic alternate routing

b) Recovery from interruptions to service, both transient and catastrophic

3.4.2 Objectives

To establish a set of accountability criteria which will provide protection against loss of data due to:

1) Transmission failures either transient or permanent

2) Failures of the DAX to which the calling subscriber is connected (originating DAX)

3) Failures of a DAX carrying tandem traffic

4) Failures of the DAX to which the called subscriber is connected (terminating DAX)

3.4.3 Approach and Progress

3.4.3.1 Accountability of Class I Traffic

Several types of accountability will be maintained for Class I traffic:

a) Overall call accountability
b) CCIS Message Accountability
c) Routing Accountability
d) Channel Accountability
e) Precedence Accountability
f) Security Accountability
g) Trunk Group Accountability (Traffic Statistics)

These are discussed below.

3.4.3.1.1 Overall Call Accountability

Overall call accountability can be maintained by what is commonly known as an Automatic Message Accounting (AMA) Function. A history of each call would be recorded containing the following information:

$$\left.\begin{array}{l}\text{Time of Call Initiation}\\ \text{Calling Party Identification}\\ \text{Called Party Identification}\\ \text{Precedence Level}\\ \text{Security Classification}\end{array}\right\}\quad\text{Block 1}$$

Time call was answered - Block 2
Time call was released - Block 3

The block containing the first five items would be recorded at the time the call is initiated at the originating DAX. The second block (sixth item) would be recorded when the called subscriber went off-hook to answer the call. The third block (seventh item) would be recorded when the call is released, due to either subscriber going on-hook or the call being pre-empted.

AMA data can be recorded at a central switch handling data for the entire network, at the nearest regional switch, or even locally at each originating DAX. If it is not recorded locally the data is sent via quasi-associated CCIS messages to the appropriate regional or central DAX where it is recorded on magnetic tape or disk.

3.4.3.1.2 CCIS Message Accountability

CCIS Message Accountability is maintained by ADCCP packet accountability procedures as described in paragraph 3.4.3.2.3. In addition, CCIS messages affecting the frame map require a response and acknowledge procedure as described in Section 3.1. No recording requirement is forseen for this type of CCIS message.

3.4.3.1.3 Routing Accountability

Routing accountability for a Class I call is maintained either until the call reaches the called subscriber's home DAX or until all of the potential routes (as defined by the deterministic routing tables) have been searched and the call is blocked. No recording requirement is foreseen for Class I call routing although the traffic statistics may include a total count of blocked calls at each DAX (ref. paragraph 3.4.3.1.7).

3.4.3.1.4 Channel Accountability

Although the data contained in the switched channels is neither recorded nor checked for accuracy, accountability of the Class I region maps is maintained by the DAXs at both ends of every link as described in Section 3.3. This confirms both the number of bits allocated to each channel and the location of the channel in the frame. Detection of errors will result in appropriate recovery procedures (Ref. Section 3.5).

3.4.3.1.5 Precedence Accountability

Precedence violation will be prevented as described in paragraph 3.1.4.4 calls attempted at a precedence higher than the maximum level allowed for the calling subscriber (defined by classmark) will be refused by the originating DAX, and an error tone returned to the subscriber. Recording of such events, although possible, is not forseen.

3.4.3.1.6 Security Accountability

Security violations will be prevented as described in paragraph 3.1.4.4 Calls attempted at a security classification higher then either the calling or called subscriber will be refused at the originating or terminating DAX, and an error tone returned to the calling subscriber. Requirements for recording such events have not yet been determined.

3.4.3.1.7 Trunk Group Accountability (Traffic Statistics)

Traffic usage statistics, data collection and printout are common features of a conventional telephone network. Traffic meters are implemented in the processor memory at each switch. Information is recorded on a predetermined schedule, on demand, or when a meter overflows. The types of statistics generally include summary information for each trunk group (link) such as call attempts, completed calls, link blocking, trunk group usage in CCS (hundred call-seconds), etc. The traffic statistics may either be printed or recorded locally or forwarded to a central switch via CCIS messages.

3.4.3.2 Accountability of Class II Traffic

Several types of accountability will be considered for Class II traffic.

a) Overall Call Accountability
b) Message Accountability
c) Packet Accountability
d) Precedence and Security Accountability

These are discussed below.

3.4.3.2.1 Overall Call Accountability

It is proposed to maintain overall call accountability by means of journal entries at both the originating and terminating DAX's so that buffer allocations may be reestablished by the appropriate receovery procedures (Section 3.5) in the event of a failure at either DAX. In the case of a call addressed to more than one subscriber, a journal entry will be made at each terminating DAX.

3.4.3.2.2 Message Accountability

It is proposed to maintain message accountability by means of reference entries at both the originating and terminating DAXs. This will increase the cross-network delay by the cumulative amount of time to record the message at both DAX's. Message accountability at tandem DAXs is not feasible in view of the proposed dynamic alternate routing proposed for Class II traffic (Ref. Section 3.2) because the various packets of a multiple-packet message are each routed through the network individually and no tandem DAX can be guaranteed to handle the entire message.

It is conceivable, however that various routes of a message could be traced by means of packet accountability as described below. Message accountability will prevent loss of any part of a message due to failure of any DAX through the use of appropriate recovery procedures (Ref. Section 3.5).

3.4.3.2.3 Packet Accountability

Dynamic accountability of packets across individual links will be provided by the ADCCP packet acknowledgement procedures which permit recovery from transient or permanent link transmission failures. When used for this purpose these procedures do not add to the cross-network delay since the packets need not be transferred to mass memory as long as there is room for them in the queue within the modulus of the link sequence numbers. In a high error environment the time to record and retrieve the packet from mass storage in order to attempt a retransmission must be considered. Consideration should be given to journal recording the header block of every packet if the message tracing capability mentioned in the paragraph above is required.

3.4.3.2.4 Precedence and Security Accountability

Precedence and Security violations for Class II calls will be prevented in the same manner as Class I traffic (Ref. Paragraphs 3.4.3.1.5, 3.4.3.1.6). Each individual packet will be routed at the precedence level at which the call was initiated.

3.4.3.2.5 Message Storage at Terminating Switches

The following procedures are proposed concerning storage of Class II messages at the terminating switch in the event that the messages are not deliverable when received.

3.4.3.2.5.1 Interactive Messages or Queries

The approach proposed in Paragraph 3.2.3.7 obviates the need for long term storage at the terminating switches since the messages are not accepted into the network until it has been determined that the called subscriber is available.

3.4.3.2.5.2 Data Base Updates

It is proposed to handle data base updates in the same manner (store-and-squirt) as interactive and query/response messages. This will help to eliminate the problem of synchronizing a new data base update with an earlier data base update (of lower precedence) which is still stored in the network pending delivery.

3.4.3.2.5.3 Narrative Messages

If narrative messages cannot be delivered immediately, the terminating switch will store them for later delivery. No advisory message will be sent to the originator.

3.4.3.2.5.4 Bulk Data Transfers

Whether or not bulk data will be stored at the terminating switch will be dependant on the size of the transfer. Medium size data blocks will be stored for delayed delivery with no advisory messages to the originator. Extremely large data blocks will not be stored, but an advisory message will be returned to the subscriber.

3.5 RECOVERY PROCEDURES

3.5.1 Problem

The DAX Communication switching network must be designed with a high degree of system availability and maintainability as an integral part of the overall plan. This is necessary to insure that traffic can be provided continuous high quality service without unnecessary delay. To accomplish this goal requires good error protection, service with a minimum of interruptions, fast restoral of service, and minimum loss of information when interruptions do occur.

This report discusses the hardware and software implications for detecting and sectionalizing troubles in the DAX network, and the recovery procedures used to restore service quickly. The impact of these recovery techniques, for transient as well as catastrophic interruptions, will be investigated as a function of traffic class.

3.5.2 Objectives

1. Identification of the hardware & software features necessary to detect and isolate faults.
2. Recovery procedures which provide a minimum disruption of service and preserve message information while testing & configuring the system around faulty units.
3. Consistency with the overall reliability, availability, and servivability goals of the network.

3.5.3 Approach and Progress

3.5.3.1 General

A successful military communication system is measured by its ability to provide continuous, economical, and accurate service in fulfilling the communications mission, without causing unreasonable delays. To meet these objectives, as part of the overall design, the system must exhibit a high degree of reliability, availability, and maintainability. The availability of continuous high quality service throughout the design life of the system is of vital importance for customer satisfaction. Typical objectives are system downtime not to exceed 2

hours over a 40-year life span and calls handled incorrectly not to exceed .02 percent. In order to provide this service capability, a system must contend with such things as traffic overloading, transmission - related phenomena (e.g., error bursts), power anomalies and failures, sabotage, and other factors which tend to degrade performance.

Built-in reliability in itself is insufficient to insure that a system will exhibit this degree of availability, particularly when equipment must be left unattended for extended periods of time. Consequently, maintainability - the ability to detect, diagnose, and correct failures - must be integrated into the overall system design. Maintenance features provided would include in-service performance monitoring, protection switchover, redundancy, alarms, and the capability of rapid fault isolation and repair.

3.5.3.2 Types of Failure Encountered

We will be concerned with two types of failure - transient and catastrophic. A *catastrophic failure refers* to problems which deny some portion of the DAX network the ability to provide basic communication services to subscribers . Such would be the case when a transmission link becomes unavailable or a DAX is not functioning due to enemy action or power failure. On the other hand, transient failure is concerned with problems which restrict the ability of the network to provide effective high quality communications for a limited time but do not necessarily eliminate it altogether. Examples of this would be loss of synchronization due to error bursts or software malfunction due to processor error. Table 3-1 lists examples of typical failures encountered under each category. In each instance the procedure will be to identify the particular problem, determine the probable cause, and take remedial action.

Survivability of the DAX network refers to the ability to satisfy the basic communication needs of surviving terminals after disruptive forces have occurred on the network. These forces may be the result of enemy action, disruptions in the transmission environment, or equipment breakdown. The survivability of a network is dependent on its topology, e.g., are nodes located far enough apart so that two or more nodes cannot be easily destroyed, are offices multi-homed, is there a diversity of path connections? Although these are important considerations, we will only be concerned here with the ability to rapidly reestablish

TABLE 3-1. Types of Failure Encountered

Catastrophic	
	Power Failure
	Extreme Traffic Congestion
	Enemy Action
	Transmission Facility Failure
	Control Transfer Failure
Transient	
	Processor Bit or Word Error
	Traffic Demand
	Transmission Related Error - Fading, Error Bursts, Impulse Noise, Phase Jitter, Envelope Delay
	Loss of Synchronization
	Environmental Failure - Lightning hits, Storm Damage

the route after disruptions have occurred. A more detailed investigation of survivability will be the subject of a future report.

3.5.3.3 Maintenance Plan

Maintainability of the switch should be incorporated as an integral part of the overall configuration. The maintenance goal is to recover from faults before service is appreciably affected so that the user is unaware of trouble. To accomplish this goal, errors and faults must be detected quickly before incorrect information propagates through the system making detection that much harder. Diagnostic software and reconfigurable hardware are the means by which to do this.

Continuous hardware checks provide for detection of faults in the pooled circuits and processing units. When a fault occurs, it is automatically detected and an interrupt in the central processor causes program control to be transferred to a series of fault recognition programs. These programs attempt to

isolate and localize the faulty unit and provide for automatic swithover of the functional operations of the perturbed and/or failed element to an identical operational element. The standby duplicate unit is immediately available for use since it has been run in synchronism with the active unit to keep its contents up to date. Thus redundancy and switchover are protective features which make the system less vulnerable to failures in individual units, hence increasing uptime by insuring that the system will operate in the presence of troubles. In a similar manner, an alternate protection channel or path diversity is provided for in the event of a valid signal loss or extreme fading.

Diagnostic software is separated into on-line and off-line programs. On-line programs are resident and operate concurrently with the operational program. They are used for recognizing that a fault has occurred and restoring system capability as quickly as possible to avoid destroying information that is currently being transmitted. The programs will cause removal of a faulty unit from service, if necessary, however, in some cases no action need be taken other than to record that an error has occurred. An attempt is made to resume normal processing at the point where the fault occurred in as smooth a manner as possible.

Off-line programs are stored in cassettes/discs and are manually loaded into the processors as required by the maintenance or technical control operators. They are used when the on-line fault recognition programs experience difficulty in restoring service. They consist of analysis and diagnostic routines which exhaustively test suspect units and record system interrupts and troubles in an attempt to isolate units in order to locate faults.

Table 3-2 lists the major phases of the maintenance philosophy used to assure system availability. Figure 3-14 shows the task assignments and allocation of equipment at a typical maintenance position. Our concern at this point is mainly with respect to recovery procedure. When irregular operation is detected, the control subsystem must react in such a way that will verify the fault, eliminate the faulty hardware, and restore the system to fault free operation.

In summary the maintenance plan exhibits the following characteristics:

a) Redundant units are used to provide service in the presence of failures and for routine preventive maintenance.

TABLE 3-2. MAINTENANCE METHODOLOGY

Establishment of a Failure Occurrance

- Detection of the Failure
- Verification that a Failure has occurred
- Classification of the Failure as Catastrophic or Transient
- Fault Localization to a Point or Points in the Network and Communication to Central Node
- Indication of Points in the Network where Fault had an Impact

Reconfiguration to Restore Full Operation

- Disabling the Unit in which the Existence of a Failure has been Established
- Removal of Failed Unit
- Replacement with a Spare Unit (or Switchover)
- Initialization of Replacement
- Resumption of Normal Operation

Repair of the Failed Element

- Localization of Failure Indications
- Isolation of a Unit which is replaceable
- Replacement and Power-Up
- Initialization Checkout
- Restoration and/or Recovery
- Enablement of Switch-In (if applicable)

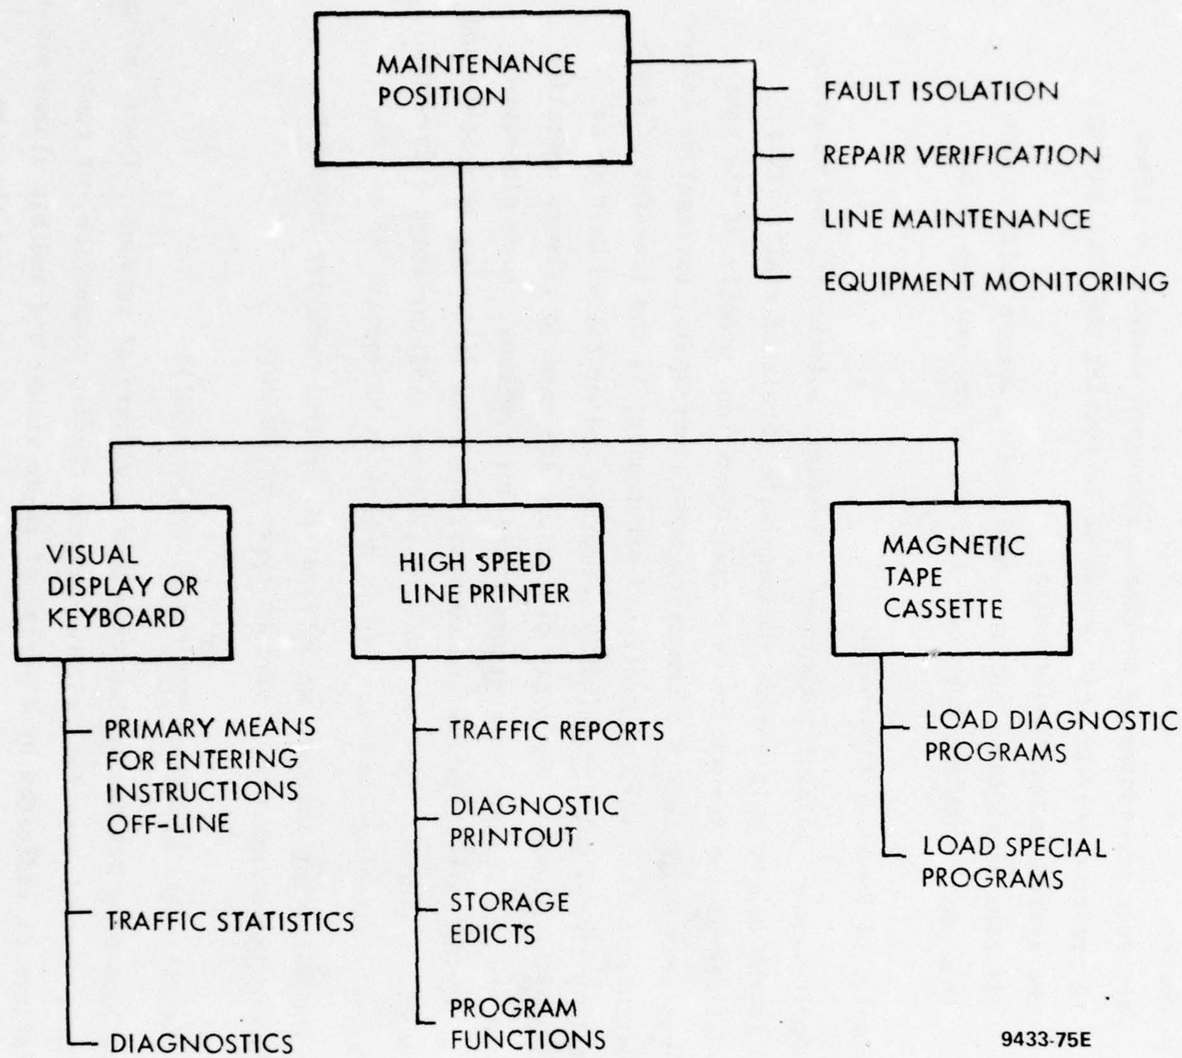

Figure 3-14. Task Assignments and Allocation of Equipment at Maintenance Position

b) Rapid detection of faults by continuous hardware and software checks.

c) Recovery procedures by on-line maintenance software designed to preserve message information while testing and configuring the system around faulty units.

d) Off-line maintenance software designed to detect malfunctions in a unit which has been switched out of the on-line system.

3.5.3.4 Examples of Recovery Procedures

A combination of several techniques - Restart, Switchover, and Recovery - is employed in the DAX to help ensure the requisite levels of availability. Restart is defined as the resumption of normal operation, usually of the same unit which has been subjected to a temporary non-catastrophic, external or internal perturbation. Switchover, as mentioned previously, is the transfer of functional operation from a catastrophically perturbed and/or failed unit to an identical operation element. Recovery procedures are used to restore operation after failure and use restart to overcome transient problems. More elaborate recovery, involving initialization and restart, are used following switchover and for catastrophic failures. Typical software programs used for store & forward, circuit & packet switching are described in detail in References 43 and 66.

In the following examples we outline the general recovery procedure followed during the occurrence of various types of failure.

3.5.3.4.1 Catastrophic Failure - DAX Cannot Process Calls

The loss of a DAX may be the result of a variety of reasons. These include power failure, control transfer failure, extreme traffic congestion, or enemy action. Failure is indicated by a series of major visual and audible alarms which alert maintenance personnel as to what has happened. In the event that the alarm at the switch is not detected promptly, the failure may be detected by congestion at other switches.

When a switch has failed or can no longer process calls, off-line maintenance programs must be used. The four primary functions of the off-line maintenance process are detection, isolation, repair, and verification. Since detection has been accomplished by alarms, the next step is to identify the probable cause of the failure. There are a variety of causes which could produce the condition that was identified. Hence it is necessary to eliminate those causes which could not produce the fault data. This process is then extended to those causes which could not influence or contribute to the fault data. Finally, our consideration is reduced to a small number of causes from which a probable cause can be chosen (subject to verification). The objective is to isolate faults to a single replaceable unit and in most cases to a replaceable cords. After a faulty unit or card has been isolated, it can be removed and replaced manually and the good unit can be verified. Verification is accomplished with monitor, test, and supervisory programs, some of which may also have been used in the isolation portion of the process.

While all this is going on System Control (SYSCON) could be reconfiguring the network to patch around the faulty DAX and operate as if that node had never been present. Subscribers originally serviced by the failed DAX would be serviced by alternates. It is assumed that multihoming of subscribers is prevalent so that a simple change in classmark is necessary to accomplish this. Class I calls in progress at the time of the outage are assumed to be lost unless provision is made to fail "softly". This would imply independence of the memory portion of the DAX used for call connection from the main processor. Under this condition connections would be frozen during the outage interval. No calls would be lost; however, no new calls could be initiated to these parties unless the calls were broken down. This last problem could also be overcome if the supervision of calls for on-hook and an off-hook indication and origination and termination of calls were done in a microprocessor which also operated independent of the main processor. Thus it is possible to conceive of a situation where none of the Class I calls are lost during an outage.

Class II calls are relatively unaffected by a DAX outage. DAX's receiving Class II packets from a failed DAX can assume that the rest of the message is lost and pass the information forward along the route. It is then the responsibility

of the switch preceeding the failed DAX to route the remaining packets through the portion of the network not affected by the outage. Packets lost at the time of the outage can be identified by the terminating switch and resent via ARQ. DAX's transmitting Class II packets to a DAX which suddenly fails must alternate route these packets around the outage to their destination. In this case, packets for which acknowledgements have not been received must be retransmitted.

3.5.3.4.2 Transient Failure - Loss of Master Frame Synchronization

Loss of frame synchronization can occur due to a transient error in the Frame Maintenance Unit (FMU) or under extreme fading conditions. When framing bit errors exceed a specified threshold, a status alarm will indicate the out-of-sync state to the control processor. This causes a Request-For-Frame pattern to be transmitted from the DAX which is out-of-sync which notifies adjacent DAX's that sync between it and the affected DAX should be reinitiated. The processor then commands the FMU to initiate a frame search and transmits an out-of-sync pattern consisting, for example, of a sequence of N-bit codes in the framing bit positions and "zeroes" in all the other bit positions*. If necessary, framing bits are sent in the return direction also. When synchronization is achieved normal data flow is resumed and an in-synchronization message is sent to the processor status indicator.

3.5.3.4.3 Transient Failure - Processor Word or Bit Error

In this example we assume that a word or bit error has affected the Class I memory map in the processor. The DAX could become aware of this situation by (a) the Class II region not starting where it should, (b) an inconsistency brought out by a periodic verification message, or (c) notification by Tech Control. Under this condition one DAX is designated master control and an algorithm is invoked where the first step consists of a message sent to validate the allocation of Class I calls. In Figure 3-15, the error has occurred somewhere within the range of bits allocated to channel j; therefore, the memory map will be verified up until the point corresponding to the start of this channel. Call allocations

---

*See Frame Synchronization (Section 4.2) for more detail.

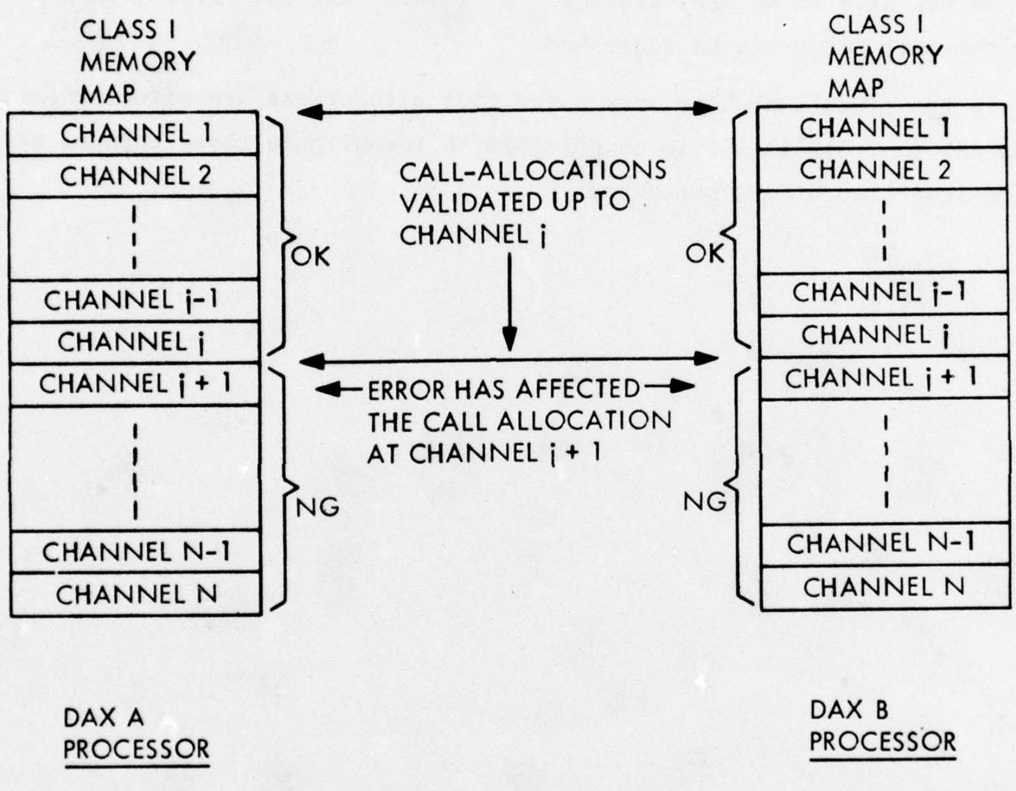

Figure 3-15. DAX Class I Map During Processor Word or Bit Error

prior to this (channels 1 to j) are unaffected by the error and are left alone. Call allocations after this point (channel j+1) are reconfigured to rapidly reconnect the calls which are affected in order to keep service loss at a minimum. It is also possible that the Class II region is affected by the error in the Class I region; in this case the Class II region would be disturbed. If it is not possible to identify in which DAX, A or B, the error has occurred, an off-line memory map may have to be made available to insure that the correct information used in the call setup can be addressed.

In the situation where only a few call allocations are affected (e.g., channels J+1 to J+3), it should be possible to reconfigure these channel assignments and leave the others untouched.

SECTION 4

SIGNALING, SYNCHRONIZATION, AND ERROR CONTROL

SECTION 4

SIGNALLING, SYNCHRONIZATION, AND ERROR CONTROL

4.1 SUBSCRIBER SIGNALLING

4.1.1 Problem

The DAX must be capable of providing service to a wide variety of subscriber terminals. These numerous types of terminals, however, fall into three general categories:

A. Telephone instruments - transmitting at various data rates (2400 b/s to 50 kb/s) and employing either digital or analog signalling and supervision.

B. Data terminals - transmitting at various data rates (75 b/s to 200 kb/s), using either character or message formats, various procedures, and operating either synchronously or asynchronously by character or message.

C. Computers - transmitting at various data rates (4800 b/s and above), using either message or packet formats, various procedures and operating asynchronously by message or packet.

In order to obtain service, a DAX subscriber must be capable of three phases of operation, call initiation, data transfer, and call termination. The means by which these three phases are accomplished will vary by data terminal category. Therefore, this note will serve the following purposes.

A. Identify the type of voice and data terminals with which the DAX is required to interface.

B. Identify the various signalling and supervision schemes appropriate for each subscriber category.

C. Identify the DAX interface for each subscriber category.

4.1.2 Objectives

The signalling and interface schemes discussed in this investigation, whether they currently exist, are proposed or are a combination of existing and proposed schemes, are intended to realize the following objectives:

A. Result in a short call-setup time.

B. Utilize control signals which are easily distinguishable from each other and from data, which are reliable and which have a low probability of false detection.

C. Place no limitations on subscriber bit patterns or formats, i.e., network transparency.

4.1.3 Analysis and Results

4.1.3.1 Subscriber Services

To place the problem in perspective, let us first outline the type of services available from the DAX. Basic switching services provided by the DAX include the following:

A. Class I Switched Service

- Virtual circuit switched digitized voice or data.
- Full-duplex transmission
- Synchronous time transparent operation.
- Subscriber dialed call initiation and signalling*

---

*In some cases automated dialing, or out-of-band signalling may be necessary, e.g., as with high-speed data terminals.

B. Class II Switched Service

- Packet switched service (packetizing of subscriber data performed by DAX, if necessary).
- Full-duplex transmission (Subscriber may choose to operate half-duplex or one-way.).
- Synchronous or asynchronous operation.
- Signalling performed by self-addressed message/packet headers.

C. Class I/II Switched Service

- Same as Class II except subscriber dialed call initiation.

Relating the three classes of switched service to the three categories of subscribers results in Table 4-1:

TABLE 4-1. COMPARISON OF CLASSES OF SWITCHED SERVICE WITH SUBSCRIBER RESULTS

Class of Service	Utilized by Subscriber Category
I	A, B
II	B, C
I/II	A, B

4.1.3.2 DAX Local Area Description

The DAX will be required to interface with various types of tactical and strategic equipment in its working environment. Some of these are shown in Figure 4-1. We will concern ourselves here only with those subscribers which are dedicated to a DAX. Interfacing with remote switches, central offices, etc., will be reported on in Section 10.

The local area arrangement for dedicated subscribers is shown in Figure 4-2 for the three classes of service offered. There are three component parts to the local access line: the subscriber terminal, the local loop, and the DAX interface. In the following paragraphs, each of these components is described in more detail.

4.1.3.2.1 Subscriber Terminals

Class I terminals consist of either telephone instruments or data terminals requiring circuit switched service (e.g., facsimile, mag tape equipment). Telephones transmit digitally encoded voice signals but may signal using either analog or digital techniques. Table 4-2 summarizes the signalling and supervision characteristics of the different telephones with which the DAX will interface. The CVSD class of telephone, which is currently under development, is the type projected for most DAX voice subscribers. It uses digital signalling in the form of cyclically permutable codewords. Data terminals

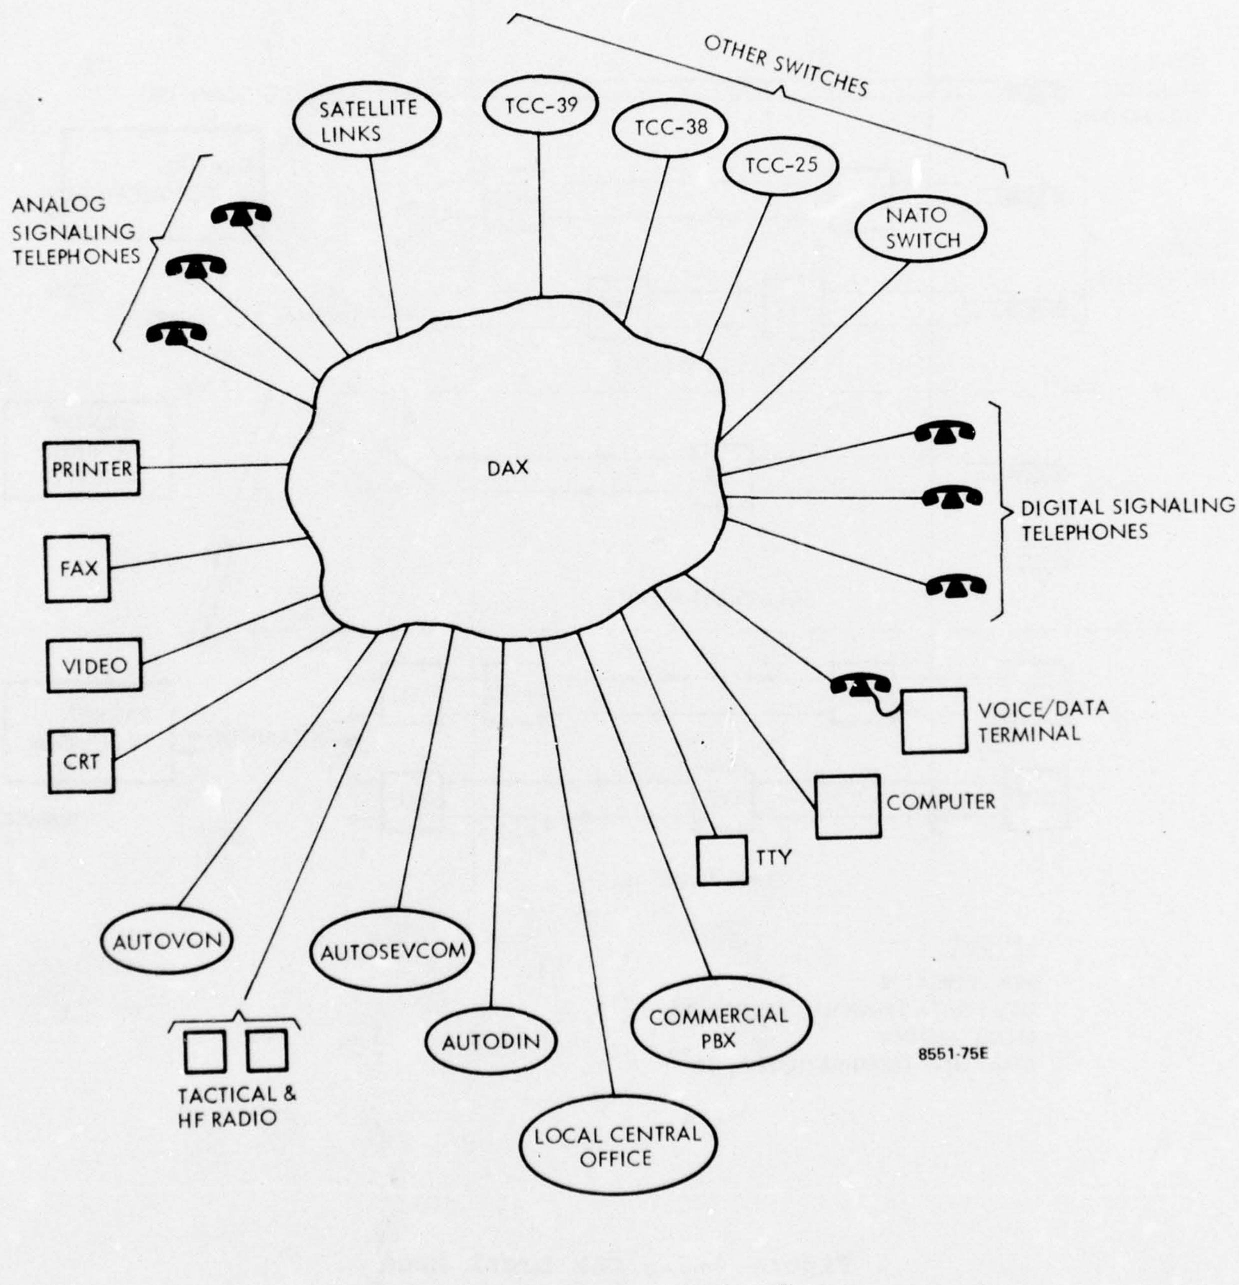

Figure 4-1. DAX Interfaces

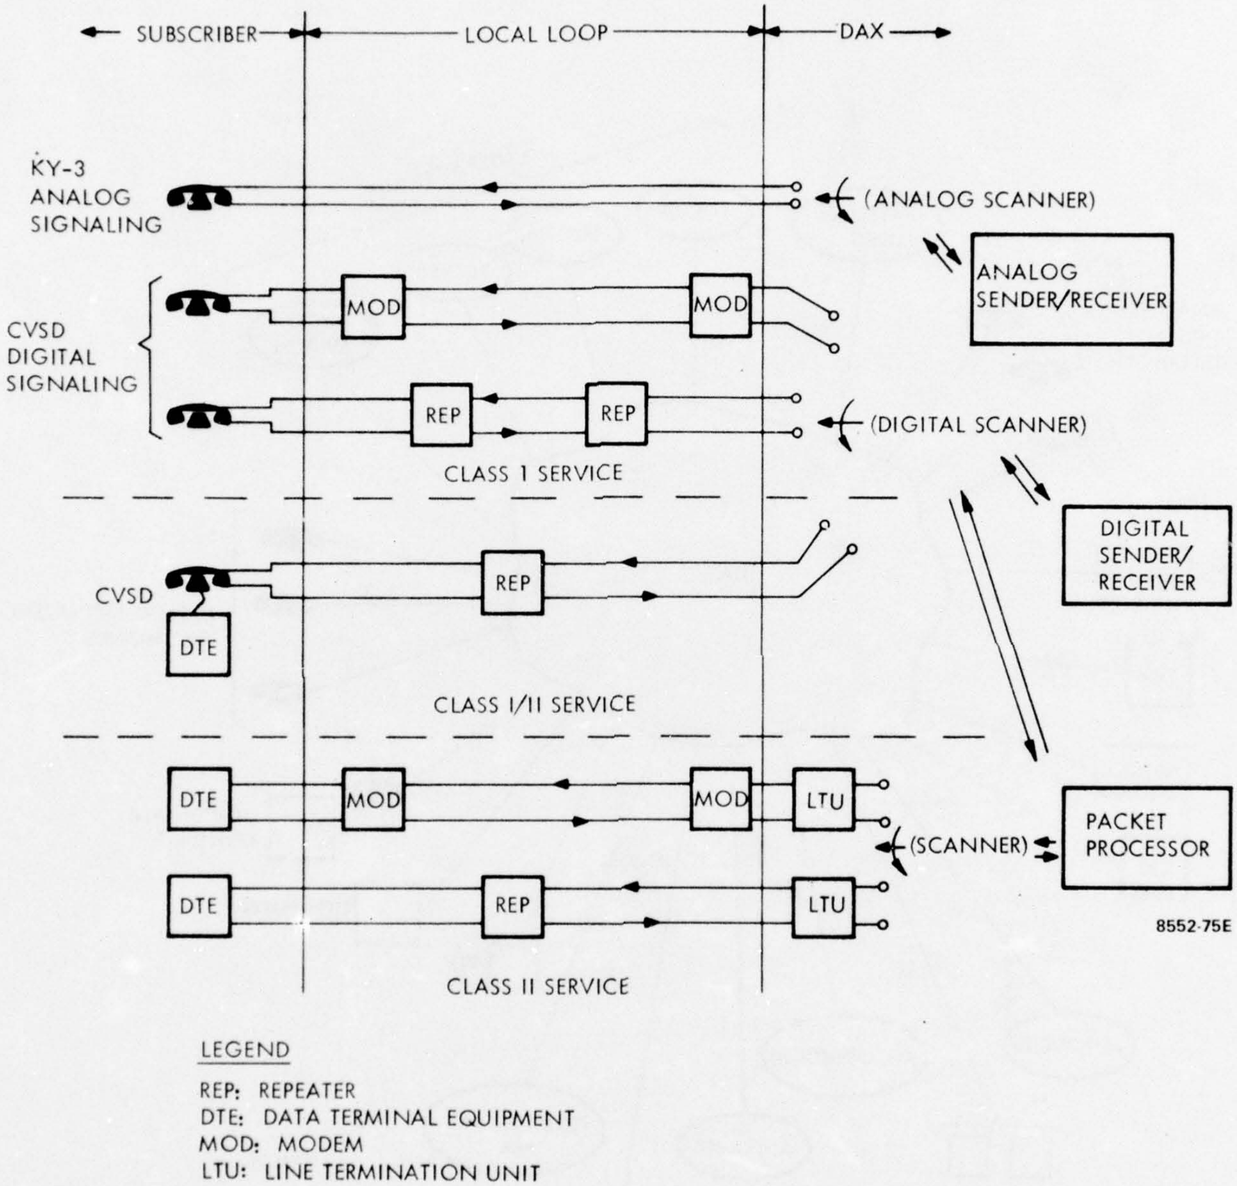

Figure 4-2. DAX Local Area

TABLE 4-2. TELEPHONE INTERFACES

l.	Interface	Supervision	Signalling
1.	TA-341	2250-Hz Seize, 2600-Hz, RLSE, 570-Hz ACK, or DC Phantom Loop Plus 570-Hz Ring	DTMF, 12 Digit
2.	TA-236	DC Common Battery Loop, 20-Hz High Level Ring	Dial Pulse, 10 pulses per sec. 60% Break, 40% Make
3.	WECO 2500 series	Same as (2)	DTMF, 12 Digit
4.	AUTOVON Telephone (local operation)	DC Common Battery Loop on transmit Pair. DC Loop Controlled Ring on Receiver pair	DTMF, 15 Digit
5.	AUTOVON Telephone (remote operation)	2600-Hz SF Seize/Release	DTMF, 15 Digit
6.	TA-720	Same as (1)	DTMF, 16 Digit
7.	TA-312	20-Hz Ringdown	20-Hz Ringdown Adapter
		Common Battery Supv Local Battery Talk 20-Hz Ring	Common Battery DTMF Terminal Ckt.
		Common Battery Supv & Talk 20-Hz Ring	Common Battery DTMF Terminal Ckt.
8.	TA-( ), CVSD	Digital	Digital
9.	KY-3	2600-Hz SF 1000-Hz Ring	SF Dial Pulse
10.	NBST	Same as (4)	Same as (4)

requiring Class I service will likely have high bit rates and consequently will use out-of-band signaling (voice coordination) or Class II type signaling (but classmarked for Class I service).

Class II terminals are composed of message and packet procedural equipment which include, for example, teletype, facsimile, paper tape, and computers. There are seven modes of operation which will be provided for Class II traffic in the DAX network. The first six are message oriented and are equivalent to the six modes used in AUTODIN. The seventh will be used by packet terminals. A definition of these modes is as follows:

A. Mode I - Duplex synchronous operation with automatic error and channel controls allowing independent and simultaneous two-way operation. Control characters are used to acknowledge receipt of valid line blocks and messages or to return error information. The terminal (or DAX) responds automatically to the control characters by continuing or stopping transmission or displaying action information to the operator.

B. Mode II - Duplex asynchronous operating allowing simultaneous two-way operation without automatic error and channel controls.

C. Mode III - Duplex synchronous operation with automatic error and channel controls utilizing one-way message transmission with the return direction used exclusively for error control and channel coordination response. Not reversible on a message basis. Control characters used in same manner as Mode I.

D. Mode IV - Unidirectional, asynchronous operation without automatic error and channel controls. Not reversible on a message basis.

E. Mode V - Duplex asynchronous operation allowing independent and simultaneous two-way transmission. Control characters are used to acknowledge receipt of message, perform limited channel coordination and display limited information to the operator.

F. Mode VI - Extended Mode I, suitable for high-speed operation over satellite, LOS or tropo links. Simultaneous, two-way operation with ARQ and/or FEC. Control messages are sent on high priority transmissions. Block sequence control and sequencing capability - no wait for validation before next transmission. Acknowledgment of delivery.

G. Mode VII - Packet mode equivalent to Mode I. Duplex synchronous operation with ARQ and simultaneous two-way operation. Supervisory packets are used to acknowledge receipt of valid data packets and for control.

It should be noted that all inter-DAX transmission of Class II data will be in packet form*. To achieve this, all DAXs will be required to convert non-packet Class II data (Mode I-VI) into packet form and to convert packet data leaving the network into a format which is compatible with the terminal to which the data is to be sent.

---

*All internal network packets will conform to ADCCP as described in Section 2-1.

Terminals in Class I/II consist of dual mode telephones which will provide either normal voice communication or data communication when used with a suitable data adapter. Signaling for voice calls will be the same as for the equivalent Class I voice call. When used in the data mode, however, the operator must first call the Class II switch using a special number and thereafter he will operate under normal Class II procedures.

4.1.3.2.2 Local Loop

The local loop is the transmission system used to connect a dedicated subscriber terminal to the DAX. A modem will be used to convert digital baseband signals outputted by a subscriber terminal into quasi-analog form suitable for transmission over the loop. A compatible modem at the DAX will convert the quasi-analog signals back into digital form prior to processing by the DAX. If the local loop uses baseband digital transmission, no modems will be necessary; however, regeneration of the digital signal at intermediate points may be required.

It should be noted that the actual loop interface seen by a data subscriber is the Line Interface Unit (LIU) of which the modem or regenerator is the loop termination portion. A detailed description of the LIU functions is given in Section 4.3.

4.1.3.2.3 DAX Interface

Basically, two types of interfaces must be provided by the DAX. One to interface voice subscribers (telephones) and the other to interface data subscribers (terminals and computers). Some of the possible local area arrangements and interfaces are illustrated in Figure 4-2. Voice subscribers

whether receiving Class I or Class I/II service will use one of the signaling schemes listed in Table 4-2. Each of the schemes listed there can be categorized as either analog or digital signaling. The functional operation of the various interfaces required for voice subscribers will be independent of the signaling form (analog or digital). Specifically, a time-shared scanner will scan each voice loop periodically for ON/OFF hook signals and a time-shared sender/receiver will be connected to all off-hook subscribers at the start of each call in order to communicate necessary supervisory signals. The mode of operation of the sender/receiver and scanner circuits, however, will depend on the signaling form. For example, analog signaling will require that scanners search for DC levels and/or tones while digital signaling will require that scanners search for valid codewords. The exact functions of the scanner and sender/receivers will not be discussed here; their technology and variations are well known and documented.

The suggested procedure for interfacing data terminals is illustrated in Figure 4-2. As may be seen, each data terminal receiving Class II service is interfaced by a Line Termination Unit (LTU) at the DAX. The functions of this device include buffering, reclocking, and possibly error checking (using parity bits for Modes I through VI and the ADCCP frame check sequence for packet data). When the LTU determines it has a complete message (character, block or packet) it sets a flag which is subsequently detected by a scanner. The scanner then notifies the packet processor to input the message. The speed of the scanner and packet processor must be such that there is no queueing of messages at the LTU. The functions of the packet processor include the following:

1. Generate in the proper code and format all handshaking signals required by data terminals.

2. Convert the messages and packets received from the data terminals into data packets suitable for transmission through the network. That is, perform all necessary code and format conversions.

3. Convert data packets intended for data terminals into the proper code and format.

A hybrid procedure may be used to interface terminals receiving Class I/II service. It is suggested that their loops be scanned in the normal Class I manner; however, when the receiver determines that they are addressing the Class II switch and not another voice subscriber, they would be connected via an internal DAX trunk to the packet processor. Thereafter, they would be treated as a normal Class II subscriber.

4.2 FRAME SYNCHRONIZATION

4.2.1 Problem

Data transmission in a digital network can be accomplished synchronously or asynchronously. When digital switching is included as a network function, it has been shown that synchronous operation is preferable[49]. In this light, it will be assumed that the DAX network is frequency synchronized and that the problem to be addressed is transmission synchronization on a link-by-link basis. More specifically, the problem of acquiring and maintaining master frame synchronization in a communication link between DAXs will be examined.

4.2.2 Objectives

The philosophy on which the DAX structure is based is that the transmission facilities of the digital network are to be utilized as efficiently as possible. To accomplish this aim, a synchronization plan should be employed which minimizes frame acquisition time. Such an objective must be integrated into a synchronization plan which also meets overall synchronization objectives. However, at this point in time, these overall objectives have not been defined. Consequently, only general synchronization procedures which meet the following short term objectives will be discussed:

a. Minimize resynchronization time
b. Realize a cost effective hardware implementation
c. Define software supervision which does not overburden the DAX processor(s)
d. Minimize transmission overhead

In Section 9.2 the specifics of these general procedures will be defined and the capability of the overall synchronization plan to meet a range of specifications will be analyzed. For purposes of immediate analysis, the synchronization specifications used for the AN/TTC-39 (see Table 4-3) will be employed.

TABLE 4-3. AN/TTC-39 SYNCHRONIZATION SPECIFICATIONS

1. Noise Model. The AN/TTC-39 shall meet the specified synchronization performance in a random error environment of 0.1% BER and in an error environment of 20% random bursts interspersed with a randomly distributed 0.1% BER. The burst rate shall be random with an equiprobable distribution between 1 Hz and 20 Hz, and the burst length shall be 5% of the time between the start of that burst and the start of the next burst.

2. Loss-of-Frame Detection. With random data on all channels, loss of frame shall be recognized within 100 ms with a minimum probability of 0.9.

3. Frame Acquisition. In-Synchronization or Frame Request framing pattern shall be acquired within 20 ms with a minimum probability of 0.9 in an error environment of 0.1% BER random errors. In the burst error environment, In-Synchronization or Frame Request framing pattern shall be acquired within 40 ms with a minimum probability of 0.9.

4.2.3 Analysis and Results

A brief review of the TDM synchronization literature available demonstrates a multiplicity of 'optimum' framing codes and procedures currently extant. The following paragraphs describe a procedure which it is felt merges some of the finer points of these myriad procedures with technology developed for the AN/TTC-39. It is not claimed that the following algorithm is optimal, only that it is very efficient in complying with the above short term objectives.

Those portions of a DAX concerned with frame synchronization are shown in Figure 4-3. The frame maintenance unit (FMU) illustrated is the hardware which realizes the chosen synchronization procedures under software control. An FMU is associated with each full-duplex inter-DAX link and has the following responsibilities.

a. Monitoring the link for loss-of-sync on the received channel.
b. Monitoring the link for notification by the distant DAX of its loss-of-sync.
c. Generation and detection of all synchronization sequences.
d. Coordinating resynchronization with the distant FMU.
e. Advising the DAX processor(s) of the sync status of its associated link.

The overall link framing process operates as follows. The first N bits of each master frame comprise the Start of Frame marker (SOF). When both transmission directions of a link are in-sync, the number of bits allocated to the SOF will be small, probably on the order of 5 to 15 bits. When either transmission direction goes out-of-sync, the appropriate FMU will increase N (implement a new and longer SOF) so that synchronization can be reacquired with great confidence within two frames. The number of bits N required in this case should not exceed 100. The attractiveness of this variable SOF approach is that it minimizes channel overhead and makes use of the powerful software capability of the DAX processor(s). Further justification of this procedure may be obtained by considering the fact that when a link goes out-of-sync, transmission of digitized voice over that link

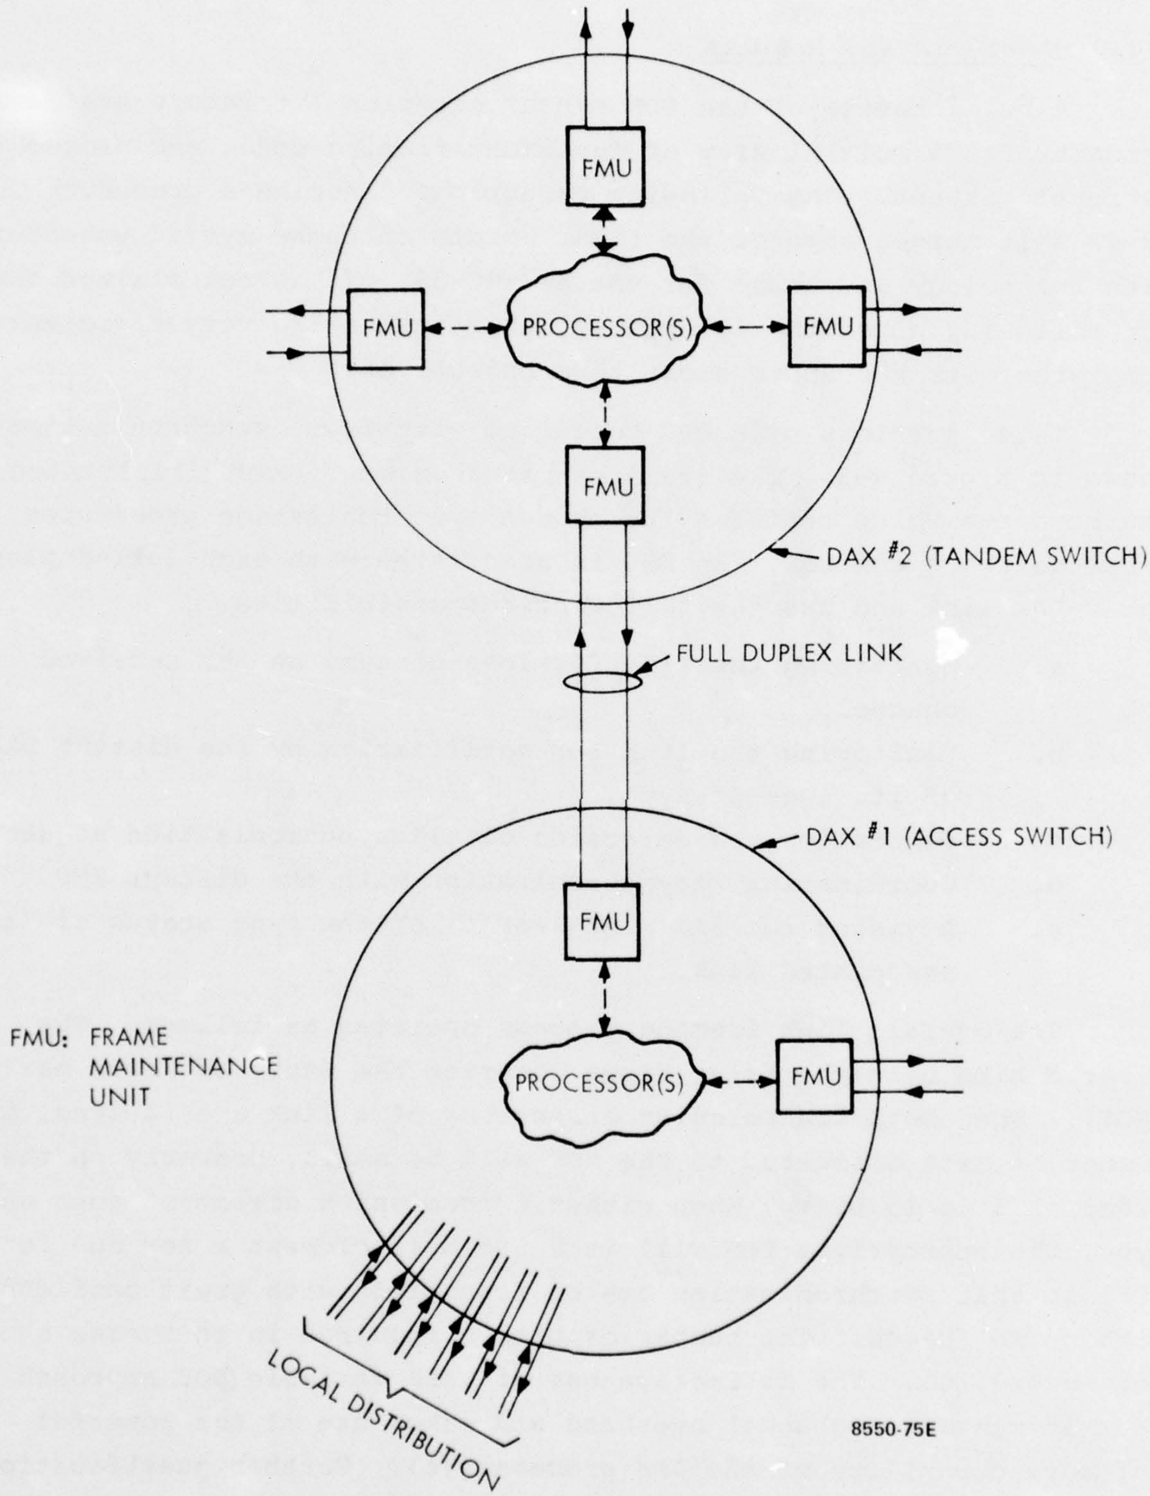

Figure 4-3. DAX Synchronization Equipment

is not possible until the link is resynchronized. Consequently, expanding the length of the SOF entails no additional inefficiency and, in fact, results in greater efficiency since it shortens frame acquisition time.

The following paragraphs provide further elaboration on specific areas of the synchronization plan.

4.2.3.1 FMU Procedures

The FMU performs two basic functions. It maintains sychronization on its receive channel and aids the distant FMU in maintaining synchronization on that FMU's receive channel. In order to accomplish these functions, three basic SOF's are proposed; N_1, N_2, and N_3. N_1 and N_2 are chosen to be of equal length and represent the short sequence mentioned in Section 4.2.3. Similarly, N_3 is the long SOF mentioned previously. The significance of these SOF's is as follows:

N_1: In-sync pattern. When transmitted from DAX 1 to DAX 2 (see Figure 4-3), it provides a synchronization pattern for DAX 2 and informs DAX 2 that DAX 1 is in-sync on the channel from DAX 2 to DAX 1;

N_2: Request-for-frame pattern. When transmitted from DAX 1 to DAX 2, it provides a synchronization pattern for DAX 2 and informs DAX 2 that DAX 1 is out-of-sync on the channel from DAX 2 to DAX 1;

N_3: Out-of-sync pattern. When transmitted from DAX 1 to DAX 2, it permits DAX 2 to resynchronize within a small number of frames.

The reason for making the length of N_1 and N_2 equal is so that they can be based on the same basic code word (e.g., complements or mirror images) and therefore have the same auto-correlation characteristic. This simplifies the hardware implementation of the FMU. With the above definitions in mind, the two basic FMU procedures required are flowcharted in Figure 4-4. Procedure 1 is used by an FMU to determine when its receive channel is out-of-sync. Procedure 2 is used by an FMU to detect an out-of-sync condition at the distant FMU. The sync and resync algorithms indicated in Figure 4-4 are taken up further in Section 9.2.

In the event that both directions of a link lose synchronization simultaneously, it will be necessary to have the capability of synchronizing on SOF N_2. An alternative or dual procedure might be to have a DAX which has lost synchronization transmit a loss-of-sync data packet (CCIS message) to the appropriate distant DAX over an alternate route in order to cause the transmission of SOF N_3. It should be noted that in this DAX scheme, loss-of-sync prohibits successful multiplexing/demultiplexing of digitized voice calls, as in any other TDM scheme; unlike these other TDM schemes, however, the DAX retains the capability to exchange CCIS messages and to continue packet data transmission.

4.2.3.2 Frame Code Sequences

For purposes of frame synchronization, a master frame is composed of a SOF and random data*. The SOF enables the demultiplexer to detect the start of the data portion of each master frame. Demultiplexing of Class I data is then accomplished by simply counting bits from the SOF with respect to previously allocated Class I bits. Class II packets are inherently self-demultiplexing.

---

* It is assumed that except for the SOF, all bits within the master frame have an equal probability of being a 1 or 0.

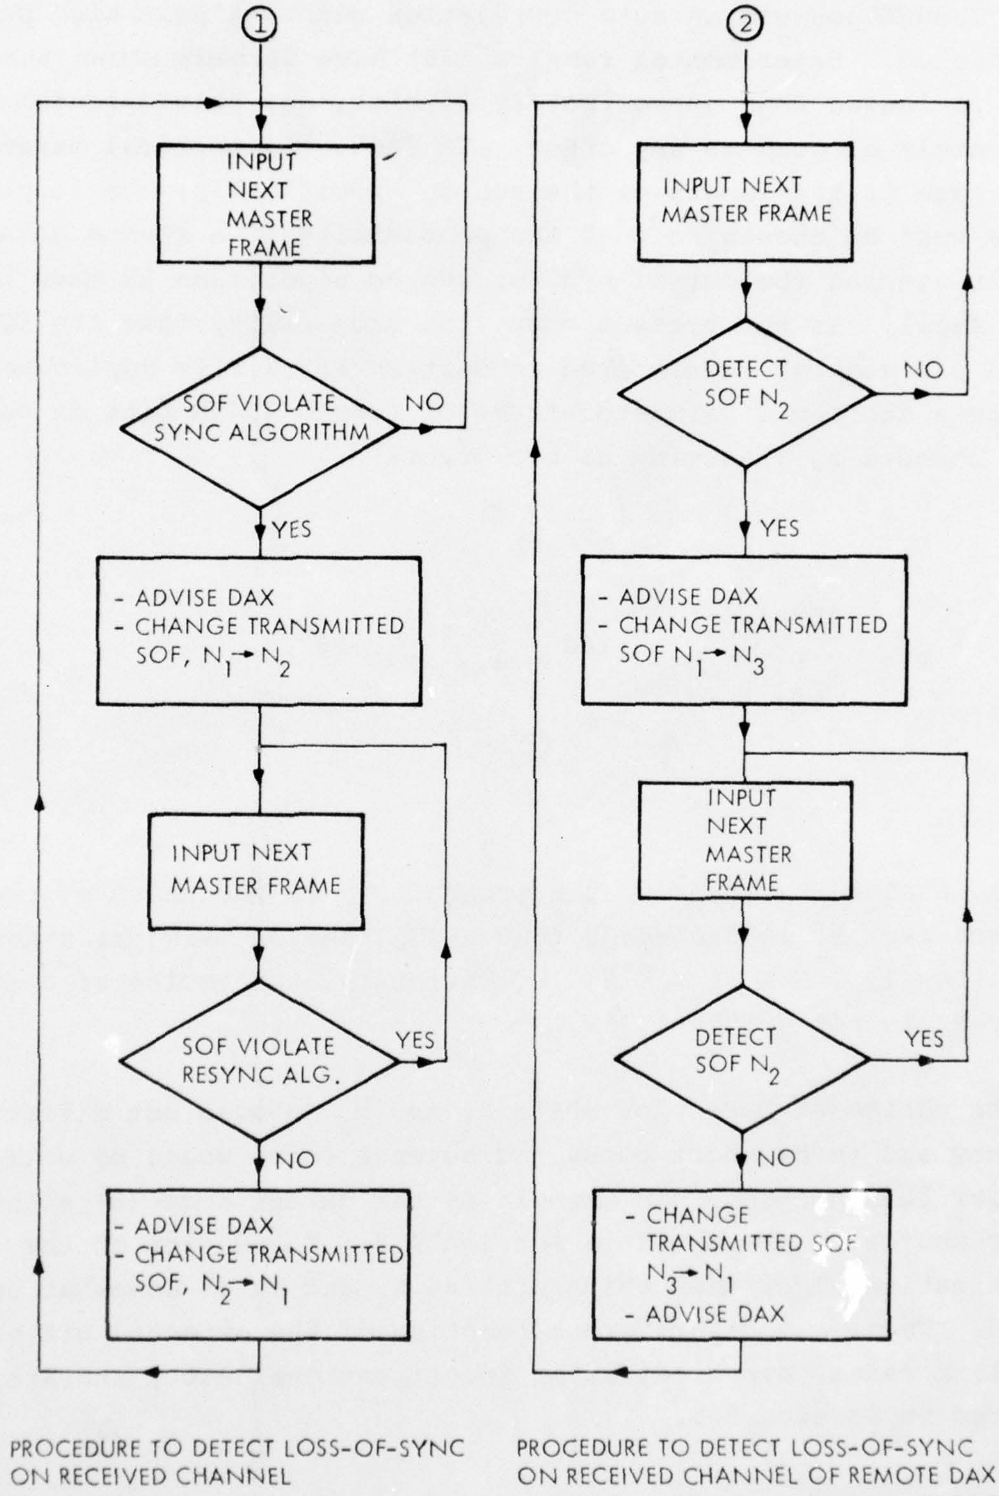

Figure 4-4. FMU Synchronization Procedures

For the type of data links being considered, it is known that the SOF should possess an auto-correlation with a single high peak and low sidelobes. Experimental results (50) have already shown that when the SOF is longer than approximately 25 bits, any pseudo-random SOF is approximately as good as any other. In fact, the critical parameter in this case is the length of the code. Specifically, the length of the code must be chosen so that the probability of a random data sequence triggering the out-of-sync or resync algorithms is made negligibly small. In the present case, the probability that the SOF N_3 (any 100 bit sequence) described in Section 4.2.3.1 is duplicated in the random data section of a master frame of length 15440 bits is conservatively bounded by (assuming no bit errors)

$$P \leq \sum_{L=1}^{15241} 2^{-100} = 1.2 \times 10^{-26}$$

which is obtained by bounding the probability of the union of events $\{S_i\}$ where each S_i is the event that a duplication sequence starts in bit position i, $101 \leq i \leq 15341$. Obviously, the choice of a specific code for N_3 is not critical.

The choice of codes for SOF's N_1 and N_2 is also not difficult since they are to be short codes and several codes would be well suited for this purpose. An example is the Barber code (6) although a better choice is presented in Section 9.2. The design of the synchronization algorithms which utilize N_1 and N_2 is somewhat complicated. These algorithms are a function of the expected bit error rates, link fades, synchronization specifications, etc., and are considered in Section 9.2.

4.2.3.3 FMU Hardware

The FMU is responsible for generating and detecting all SOFs. It is expected that detection will be accomplished using some form of the optimum correlation detector (48). Although this detector is ideally suited for the short sequence N_1 and N_2, the length of N_3 might make it impractical to incorporate an N_3 detector in each FMU. A cost-effective approach would be to time share a single N_3 detector among the FMUs under the supervision of the DAX processor(s).

4.2.3.4 Trunk Encryption

The impact of trunk encryption on the overall frame synchronization plan has been ignored for the time being. This aspect of synchronization is examined in Section 4.5, 3.-. It may be noted that the DAX can handle encrypted Class I calls easily, as can a conventional TDM scheme. However, application of modern encryption techniques will have the effect of converting a character, block, or packet-oriented data terminal into a psuedo-Class I terminal (i.e., one emitting a continuous, synchronous, binary-symmetric bit stream). At least two possibilities warrant exploration; one is to handle encrypted Class II calls as psuedo-Class I calls; the other is to treat the encrypted bit stream as pure text, and to "super-packetize" this "text" data, analogously to trunk super-encryption of a Stage II call. Additionally, a compromise is conceivable such that high speed crypto bit streams (e.g., 2.4 Kb/s through 32 Kb/s) would receive Class I handling, and low speed crypto (e.g., 45 B/S through 1200 B/S) would be converted into packets.

4.2.3.5 Synchronization Implementation

In Section 9.2 the relative cost effectivity of the short, short complement and long SOF approach is weighed against the cost effectivity of a short, short complement, and iterated short or short complement approach.

4.3 TRUNK/LOOP SYNCHRONIZATION

4.3.1 Problem

Within its local serving area, the DAX must be capable of communicating with a wide variety of subscriber terminals. This Task examines the following aspects of this communication problem:

(a) Specify appropriate means for maintaining bit integrity in the local loop, where the local loop is defined to be any transmission system connecting a dedicated subscriber to the DAX;

(b) Determine the impact of various local loop and subscriber terminal operating characteristics and control functions on bit synchronization;

(c) Define the timing and control required to transfer information bits between a local loop and an internode link by DAX class of service

4.3.2 Objectives

(a) To integrate local loop synchronization into overall network synchronization, thereby providing end-to-end bit integrity.

(b) To identify the system functions to be performed by the loop terminating device at the subscriber location.

(c) To identify the type of signals required at the interface between the subscriber terminal and the loop termination device.

4.3.3 Analysis and Results

4.3.3.1 Background

A prerequisite for reliable digital communications between a dedicated subscriber and its assigned DAX is a common time scale; that is both the DAX and the subscriber must possess local timing supplies which operate, to a high degree of accuracy, at the same frequency. All timing signals at a DAX are derived from a nodal clock. It may be assumed that the nodal clock is an extremely accurate and stable oscillator and that any timing signal required whose frequency differs from the nodal clock frequency is digitally synthesized with the same accuracy and stability as the nodal frequency. Nodal clocks maintain long-term frequency stability through a network synchronization plan as described in Section 4.4.

Subscriber timing for synchronous operation can either be locally generated or slaved to network timing. In the vast majority of cases (which includes all voice terminals and most data terminals), subscriber terminals accept timing from the network. In this mode, the subscriber terminal equipment uses clock signals provided by a line interface unit* (LIU) to time both its transmit and receive digital signals. If the terminal equipment cannot slave to external timing, as is the case with most asynchronous terminals, it will be

*A device which interfaces the local loop and the DTE.

necessary for the network to buffer all data passed between it and the subscriber in order to insure bit integrity. This buffering when necessary will be accomplished within the LIU.

It may also be necessary to establish other types of synchronization before communications can occur; these other synchronizations may require identifying characters, bit patterns (codewords), messages or packets. As an example, AUTODIN message oriented terminals require establishing both character and message synchronization prior to data transfer. However, the exact nature of these synchronizations depend upon the code, format and mode of operation used. Similarly, packet oriented terminals operate under a multiplicity of packet synchronization schemes. The particular packet protocol used in the DAX network for CCIS and Class II data requires identifying the 8-bit codeword (flag sequence) 01111110 as signaling the start of a packet. Character synchronization is not necessary here since the protocol is bit oriented (see Section 2-4). A number of other packet systems are based on ASCII characters and therefore require both character and packet synchronization.

Although these other types of synchronization are an important aspect of local loop transmission, they are basically software extensions of bit synchronization and are more correctly considered as part of the subscriber protocol. Obviously, the DAX will conform to all subscriber protocols on an individual basis as described in Section 4.1. Section 4.3.3.2 discusses in detail the subject of bit synchronizing subscriber terminals.

Another aspect of access synchronization concerns the timing and coordination required to transfer Class I information bits between local loops and internode links, so as to maintain the integrity of Class I virtual circuit switched service. In a conventional synchronous TDM/digital switching system, once a subscriber is allocated a channel in a link out of his local node, the allocation remains fixed for the duration of the call. Consequently, the timing and cross-office delay associated with that call, as established during the connection phase, also remains fixed. However, with the FTDM scheme employed in the DAX, the channel allocation of a particular call

within its master frame may change many times during the course of the call. As a result, nodal channel coordination is considerably more complicated than in the conventional TDM. Section 4.3.3.3.1 examines this subject in greater detail.

4.3.3.2 Bit Synchronization

As described in Section 4.3.3.1, it is necessary for the network an subscriber terminal to be in bit synchronism. There are two approaches for achieving this requirement, either the subscriber terminal slaves to network timing or the network, in effect, slaves to subscriber timing. In the latter case, the LIU would buffer data received from the subscriber in order to absorb clock differences.

It is assumed that the LIU provides the hardware and, if necessary, software to achieve bit synchronization. Strictly speaking though the LIU would only interface data terminal equipment; for voice service, a digital telephone would terminate the loop. As described in Section 4.1, it is assumed that the DAX will interface with all telephone signaling and supervision schemes (analog or digital) as required. To simplify matters further, bit synchronization in the local loop will only be discussed with regard to data terminals. It can be assumed that synchronization for digital telephones can be and is achieved in much the same manner.

The remainder of Section 4.3.3.2 examines a proposed LIU in order to demonstrate its required functional operations and the impact of various loop and subscriber requirements on synchronization.

4.3.3.2.1 Line Interface Unit - System Operation

A block diagram for a proposed LIU is shown in Figure 4-5. Because the device is intended to be applicable to all Class I, Class II and Class I/II data terminals, more operations are shown than might

---

*The interface circuits shown in Figure 4-5 are not being proposed as a standard DAX interface at the subscriber level; they are only intended to demonstrate the type of signals that must be exchanged.

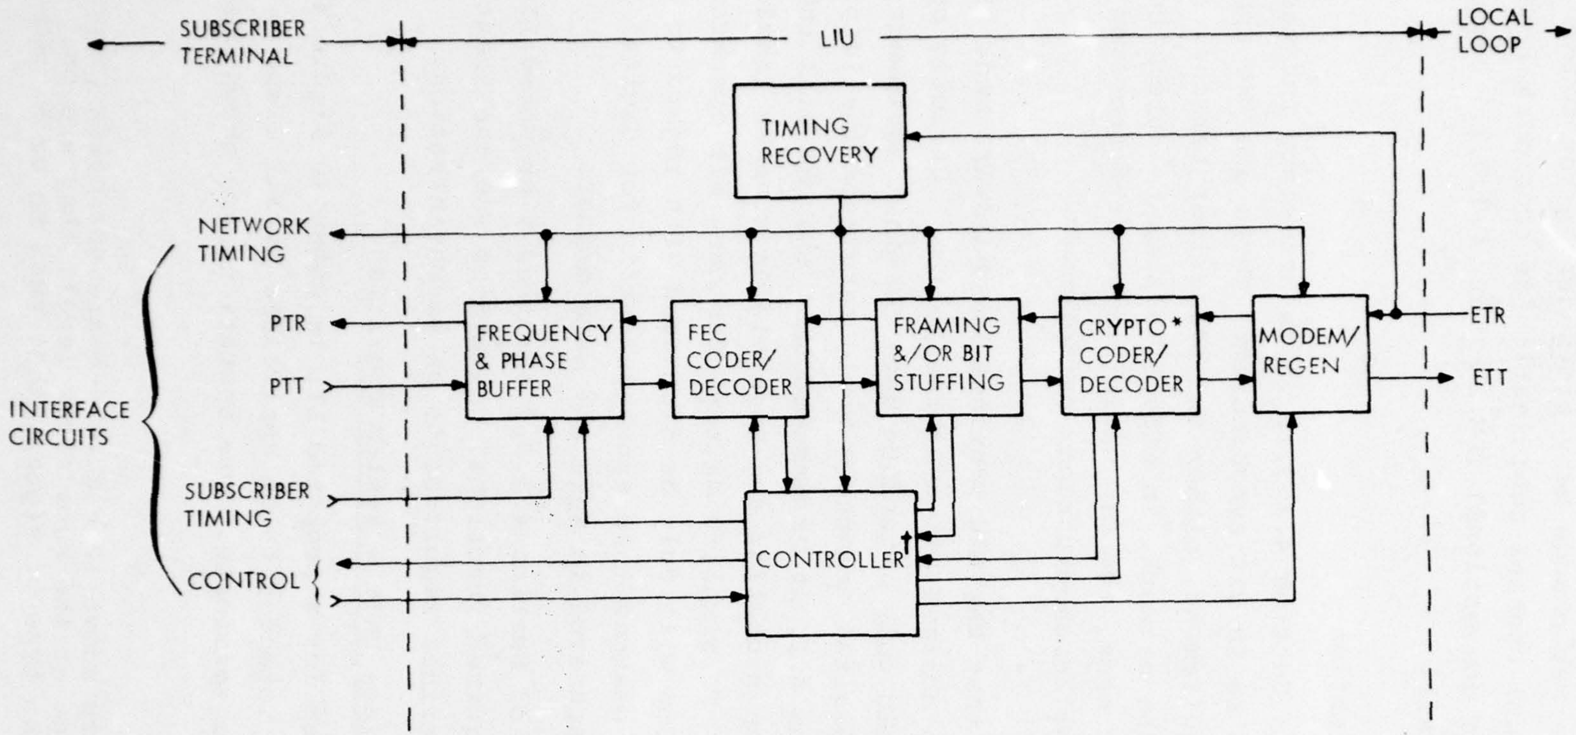

PTR: PLAIN TEXT RECEIVE
PTT: PLAIN TEXT TRANSMIT
ETR: ENCRYPTED TEXT RECEIVE
ETT: ENCRYPTED TEXT TRANSMIT

9435-75E

* GFE

† PROBABLY WILL INCLUDE A CRYPTO ANCILLARY UNIT

Figure 4-5. Line Interface Unit (LIU) Block Diagram

be required by a particular subscriber. The following subsections describe how these operations vary by subscriber and loop characteristics.

4.3.3.2.1 Modem/Regenerator

A modem will be used to convert digital baseband signals outputted by a subscriber terminal into quasi-analog form suitable for transmission over the local loop. A compatible modem at the DAX will convert the quasi-analog signals back into digital form prior to processing by the DAX. If the local loop uses baseband digital transmission, the modems will be replaced by digital regenerators.

For Class I or Class II service, the modem/regenerator will provide a transmission capacity equal to the DTE data rate. For Class I/II service, however, the modem/regenerator must be compatible with the data rate of the DTE and the digital telephone (e.g. 16 kb/s CVSD), since both are sharing the same local loop. If the DTE and telephone do not transmit at the same data rate, the modem will transmit and receive at the higher of the two rates. When the slower device is operating, the LIU will provide speed buffering.

4.3.3.2.2 Timing Recovery

Basically this will be accomplished by phase locking a local oscillator to local loop timing. This recovered clock will be the primary base for controlling LIU logic and control circuitry, data transmission and reception, and DTE timing (if the DTE is externally timed). Stability and accuracy of this time base is derived from DAX timing.

4.3.3.2.3 Crypto Synchronization

It is assumed that all local loops will be required to provide complete Communications Security (COMSEC). To this end, compatibility with encryption and description equipments have been included as LIU requirements. It is further assumed that the LIU will be capable of encrypted operation on a link-by-link basis (LIU to LKG at the DAX) or an end-to-end basis (LIU to LIU) for either local or internode calls.

Current and future crypto devices as planned by the military are bit-by-bit substitution devices. That is, a pseudo random bit stream is combined with a plain text bit stream via a reversible operation. Consequently, slipping a single bit in the loop modems will cause loss-of-crypto sync. To minimize the effects of such an occurrence, the LIU should be provided with the capability of automatically detecting and correcting loss-of-crypto sync. As an added precaution, the LIU should be able to accept a resync command from the DTE.

4.3.3.2.4 Framing/Speed Buffering

This function will likely be employed only when necessary. Examples of circumstances requiring its use are as follows:

(a) For Class I/Class II operation, speed buffering (bit stuffing) and/or framing may be necessary to match the subscriber terminal or telephone transmission rate to the modem rate.

(b) Asynchronous terminals will be made 'pseudo synchronous' by speed buffering so that a synchronous crypto device can be employed. For example, a teletypewriter signal between characters would appear as an encrypted steady MARK signal on the local loop.

(c) As a general means of detecting loss-of-crypto sync, framing all loop signals is a possibility. Loss of crypto sync would then appear to the LIU controller as loss-of-frame.

4.3.3.2.5 Forward Error Correction

It is proposed that this function be provided in noisy environments in order to increase throughput. A likely candidate for FEC is the (16, 8) quasi-cyclic code which corrects up to two errors in a 16-bit codeword which contains 8 information bits and 8 redundancy bits. The code performs well in an error environment of 0.1 percent or less and has the added advantage of accepting and outputting individual 8-bit characters, thus providing the character sync function for ASCII characters.

4.3.3.2.6 Frequency and Phase Buffer

The LIU will be capable of operating from the subscriber terminal clock when transferring data between the LIU and the subscriber terminal. However, all data transmitted from the LIU to the DAX will be timed by the LIU clock. In order to bring subscriber data into synchronism with network timing when the DTE is not timed by the LIU, a buffer will be used to read subscriber data into the LIU.

4.3.3.2.7 Controller

The controller coordinates all operations within the LIU and generates and receives all interface signals required by the subscriber terminal (e.g., Data Set Ready, DTE Ready, Resync, Disconnect, etc.). During the data transfer phase of a call the LIU will provide a transparent interface; all line protocol required of the network will be under the control of the LTU or DAX packet processor as described in Section 7. Additional functions likely to be included under controller supervision include:

(a) On-Hook/Off-Hook supervision: For Class I/II service the controller will generate all codewords required by its associated telephone in order to initiate or terminate a call. This mode of operation is analogous to that provided by the TENLEY DSVT/Data Adapter combination where the LIU would be the data adapter equivalent.

(b) Loopback: The controller will detect and initiate all loopback commands directed to the LIU by the DAX nework should this be required as part of the overall maintenance plan.

(c) Prè-emption: The controller must detect the codeword indicating a pre-empt or forced clear condition. It will then in form the subscriber terminal over the appropriate interface circuit.

4.3.3.2.8 Interface Circuits

As shown in Figure 4-5, the DTE/LIU interface consists of six functional interface circuits, one pair for exchanging timing, one

pair for exchanging clear data, and one pair for exchanging control information. Basic electrical interface circuits such as power and ground have been ignored. As mentioned previously, this is not a proposed interface, it is only being used to illustrate the system requirements placed on the LIU. At this point in time, it would not be appropriate to attempt to define a general DAX/subscriber interface which would be sufficiently flexible to handle all present and future DTE needs, especially given the uncertainties relating to the shift from character-oriented to packet-oriented data link control procedures being proposed by most Standards committees.

4.3.3.3 Local Loop/Internode Link Data Transfer

Class I virtual circuit switch data undergoes a format change at both the originating and terminating nodes due to the multiplexing and switching functions performed at these nodes. To illustrate this point consider the following: Class I data transmitted to the network traverses the local loop in the form of a continuous bitstream. At the originating node this bit stream is fragmented into sections, a masterframe period in length, consistent with the FTDM concept as described in Section 2.1. These sections of data are placed in sequence in the appropriate outgoing link, one section per allocated channel per each outgoing master frame. The channels then sequentially traverse the network and arrive at the terminating node, each subject to the same fixed cross-network delay. At this terminating node the channels are, in effect, unfragmented and transmitted over the appropriate local loop as a continuous bitstream.

As described, this fragmenting procedure is the same as would be required in any conventional TDM/digital switching system. What complicates the procedure with respect to the FTDM concept is the fact that channel allocations in each master frame may be dynamically reassigned on a link-by-link basis. In order to assure that these reassignments do not cause loss of network transparency (i.e., loss of bit integrity) all data is delayed (buffered) at each node. The extent

of the delay at tandem nodes is examined in Section 4.3.3.3.1. The extent of the delay at both the originating and terminating nodes is examined below. It should be noted that since the Class II region of the master frame provdes pseudo real time packet switched service, it is not affected by this particular timing problem.

4.3.3.3.1 Originating Node Timing

Class I data during the data transfer phase of a call arrives at the DAX in the form of a continuous bit stream. The data is fragmented in sections (a master frame period in length) in real time by the DAX and each segment is sequentially placed in its outgoing channel allocation on a frame-by-frame basis. The process is illustrated in Figure 4-6. To simplify the figure, it is assumed that there is no Class II data and that the entire master frame is composed only of variable length Class I channels. As shown, section #n is placed in master frame #n, section #n+1 is placed in master frame #n+1, and so on. Observe that in master frame #n+1, calls 1 through 3 terminate. In the subsequent master frame, the master frame contents are compacted per the FTDM concept. Due to the random temporal relationship between segments and master frames, in master frame #n+2 the data for channel 4 which is now adjacent to the start of frame marker is required before it has been completely received. This channel slip would result in a loss of bit integrity for this call since the DAX would have to repeat section #n+1 in master frame #n+2, as is shown.

To avoid this 'hiccup' effect, it is necessary to delay each section at the originating node one master frame period before placing it in its outgoing channel allocation. This assures that a channel allocation can never be reassigned (moved up in the master frame) to the extent that its data is required before it is received from the local loop. In effect, this delay aligns the master frames and sections such that each section ends as the master frame in which it is to be placed begins. This is illustrated in Figure 4-7.

Observe in Figure 4-7, that the cross office delay at the originating node is not fixed. At the start of a call, the sections will probably be placed near the end of the master frame (depending on

AD-A039 548 GTE SYLVANIA INC NEEDHAM HEIGHTS MASS ELECTRONIC SYS--ETC F/G 17/2
SENET-DAX STUDY. VOLUME 1.(U)
JUN 76 DCA100-75-C-0071
UNCLASSIFIED FR76-1-VOL-1 NL

3 OF 3
AD
A039548

END
DATE
FILMED
6-77

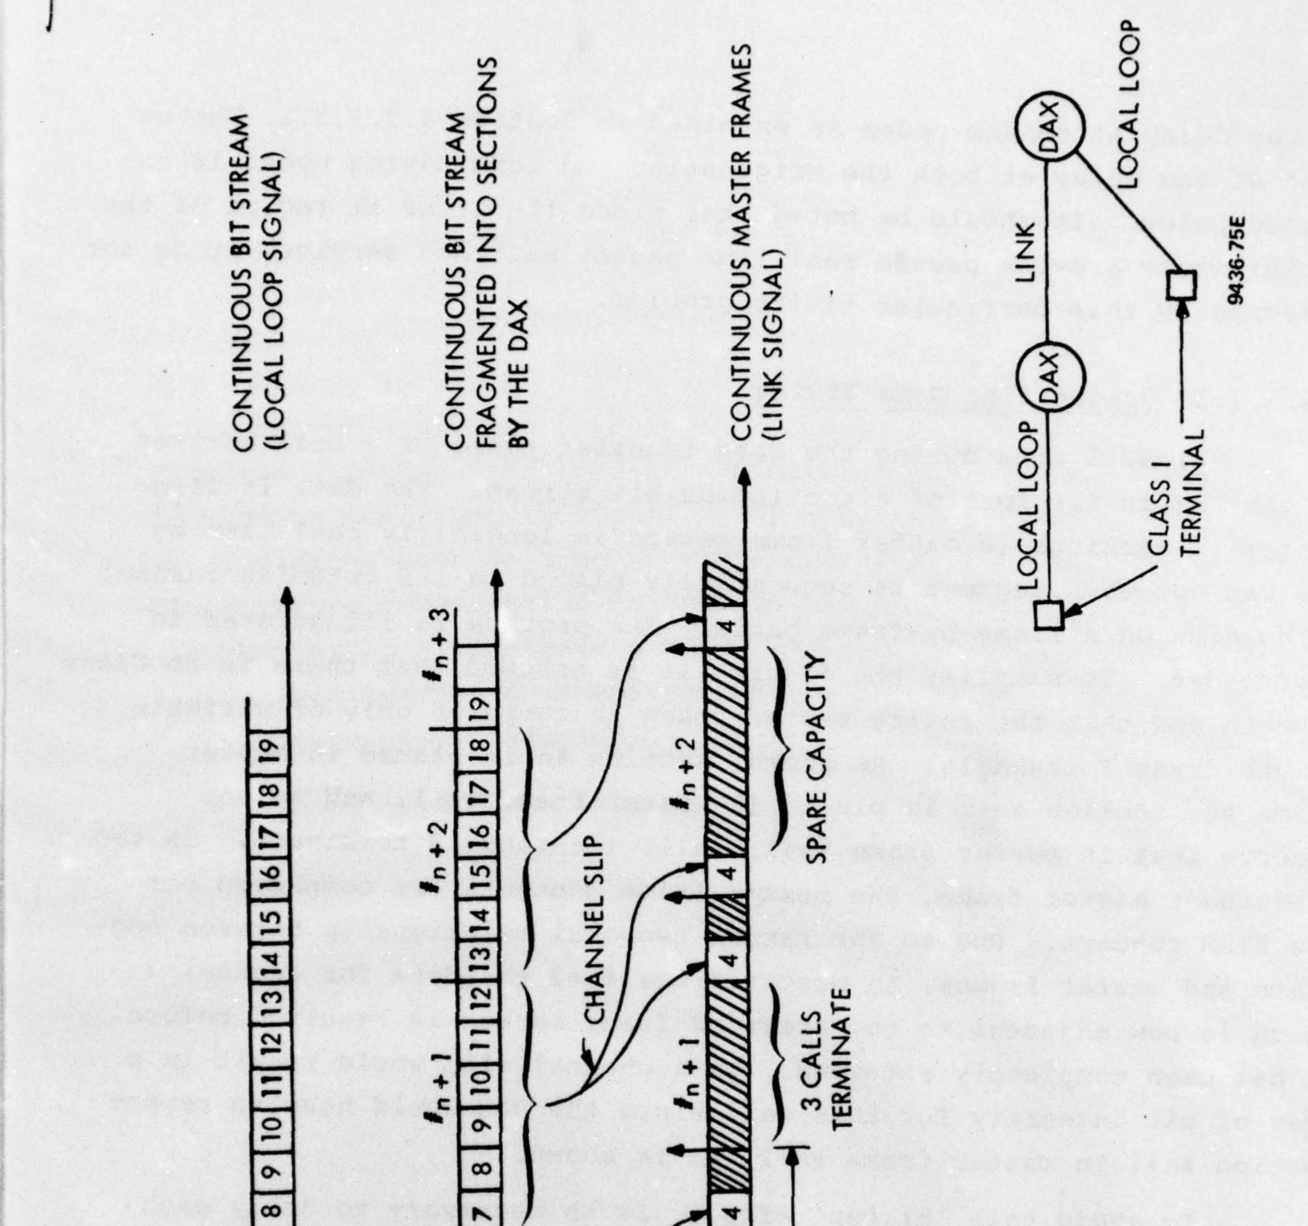

Figure 4-6. Local Loop/Link Coordination at Originating Node

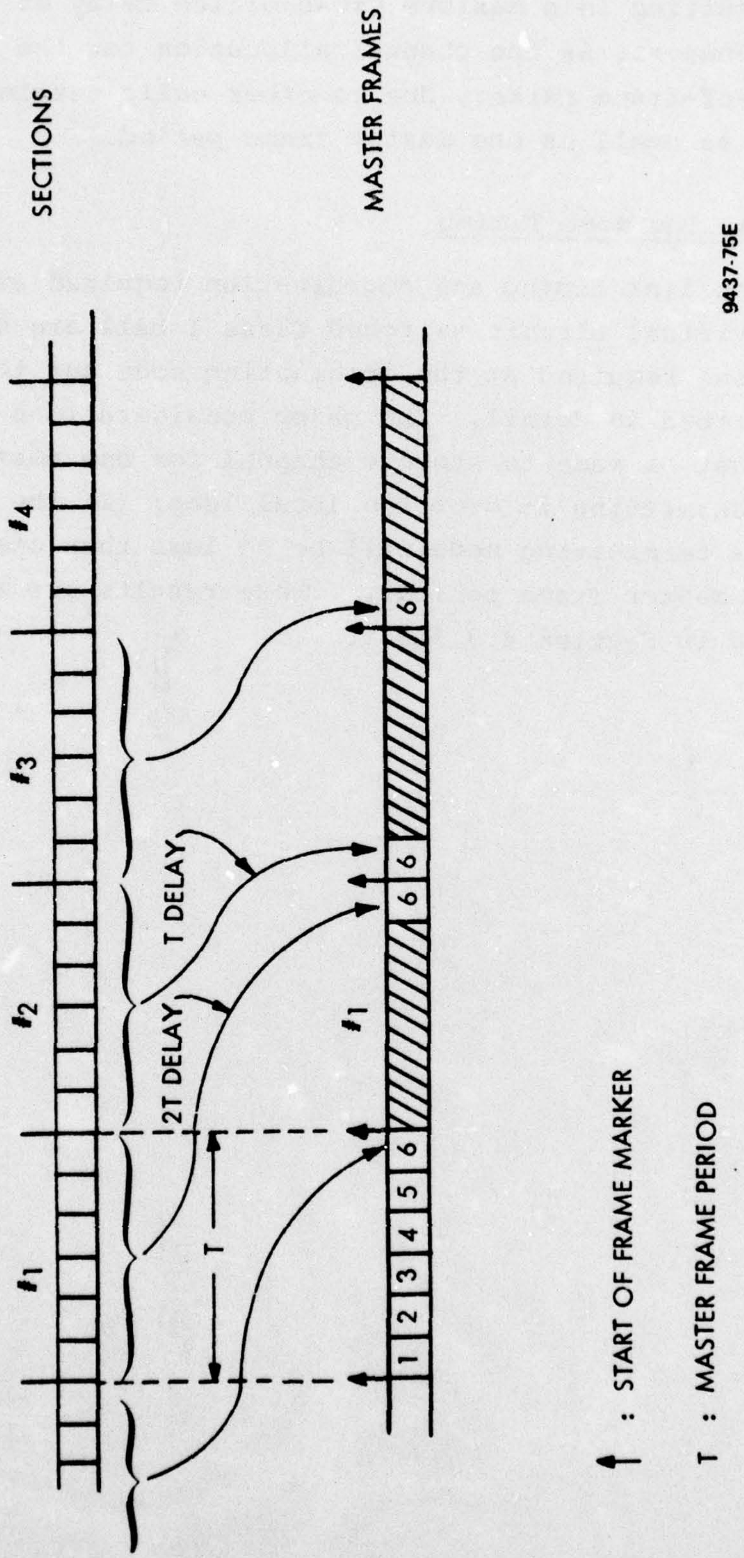

Figure 4-7. Cross-Office Delay Variations at Originating Node

traffic load) resulting in a maximum cross-office delay of two master frame periods. However, as the channel allocation for the call moves toward the start-of-frame marker, due to other calls terminating, the delay can become as small as one master frame period.

4.3.3.3.2 Terminating Node Timing

Local loop/link timing and coordination required at the terminal node of a virtual circuit switched Class I call are similar in concept to those required at the originating node and therefore will not be described in detail. The major considerations here are (1) provisions must be made to store a channel for one master frame period before transmitting it over the local loop; (2) the cross-office delay at a terminating node will be no less than one and no greater than two master frame periods. These results are analogous to those obtained in Section 4.3.3.3.1.

4.4 NETWORK SYNCHRONIZATION

4.4.1 Problem

The problem being considered is to define the means by which the DAX network synchronizes switching and multiplexing functions so that all data which is received at a node from a connected node and is to be transmitted to a connected node, is delivered to the proper destination, in the proper sequence, and without the addition of extraneous bits. The problems of synchronizing communications between a DAX and a dissimilar system such as AN/TTC-39 or a NATO switch, or between a DAX and a dedicated subscriber are examined in Sections 10.1 and 4.3, respectively.

There are two aspects to DAX network timing, internode and intranode synchronization. The first includes bit synchronization (network bit integrity) and master frame synchronization on a link-by-link basis. The second concerns master frame alignment and coordination at each node. Maintenance of bit integrity in the DAX network constitutes the same problem it does in any digital network and is discussed in detail in this section. Frame Synchronization is examined in Section 4.2 and will not be considered here. Frame alignment and coordination in the DAX is somewhat more complicated than in a convention TDM/digital switch due to the dynamic nature of the DAX master frame. This problem is considered in Section 4.4.3.2.

Throughout this discussion, it will be assumed that all links connecting nodes transmit at the same data rate and possess the same master frame period. The impact of variable link transmission rates and master frame periods on network synchronization is discussed in Section 4.4.3.5.

4.4.2 Objectives

To choose a network synchronization scheme out of the various synchronous and asynchronous plans available which is attractive with respect to the network factors of survivability, cost, reliability, and complexity.

4.4.3 Analysis and Results

4.4.3.1 Background

The network under consideration consists of a set of nodes (DAX's or similar systems) interconnected by various transmission systems (e.g., LOS radio, coaxial cable, satellite links, etc.). Because all data passing through a node undergoes synchronous multiplexing and switching, it is necessary to keep this data in frequency and phase synchronism with the local nodal clock. It is not possible however to have all nodes within a network in perfect synchronism. Consequently, any data received at a node from a connected node must be brought into synchronism with the local clock. There are two primary methods for performing this internode synchronization:

a. A synchronous or clocked approach where each nodal clock is controlled so that all clocks maintain the same average frequency.

b. An asnychronous or unclocked approach where all nodal clocks are independent, extremely stable and free running, and buffers are used to absorb frequency errors.

Within each primary method there are numerous timing schemes available. This note will specifically consider the four major internode synchronization methods and possible combinations:

a.	Independent Atomic Clocks	Asynchronous Clocks
b.	Bit Stuffing	
c.	Master/Slave	Synchronous Clocks
d.	Frequency Averaging	

In addition to internode synchronization, the necessity for intranode synchronization or coordinating the transfer of Class I channels from input master frames to output master frames at a node is also addressed. Most of the problems considered within this subject are unique in that they have no analogy in the conventional TDM/ digital switching system.

4.4.3.2 Master Frame Alignment and Coordination

In a conventional synchronous TDM/digital switching system, master frames are aligned at each node in order to simplify switching and multiplexing functions. This is illustrated in Figure 4-8. As may be seen, two sub-multiplexed channels A and B are to be multiplexed together. Each signal is first bit and frame synchronized and then frame aligned using a frame reference signal provided by the nodal timing system. The frame alignment process assures that channels 1a and 1b, 2a and 2b, etc., are in time alignment prior to multiplexing. An alternative to frame alignment is to multiplex the channels in the order in which they are received and to transmit this order via a control channel to the demultiplexer in order to insure proper demultiplexing. This alternate method is considerably more complicated than frame alignment and this complexity is the major reason it is not used.

The master frame structure employed in the DAX is shown in Figure 4-9a. For the purposes of this discussion, it is assumed that there is no Class II data and that the entire frame is composed only of variable length Class I channels. The DAX master frame is dynamic in that a channel allocation (its position in the master frame) can change from frame-to-frame; this is illustrated in Figure A-9a. Note however that after the initial allocation, a channel can only move toward the start-of-frame marker and never away from it.

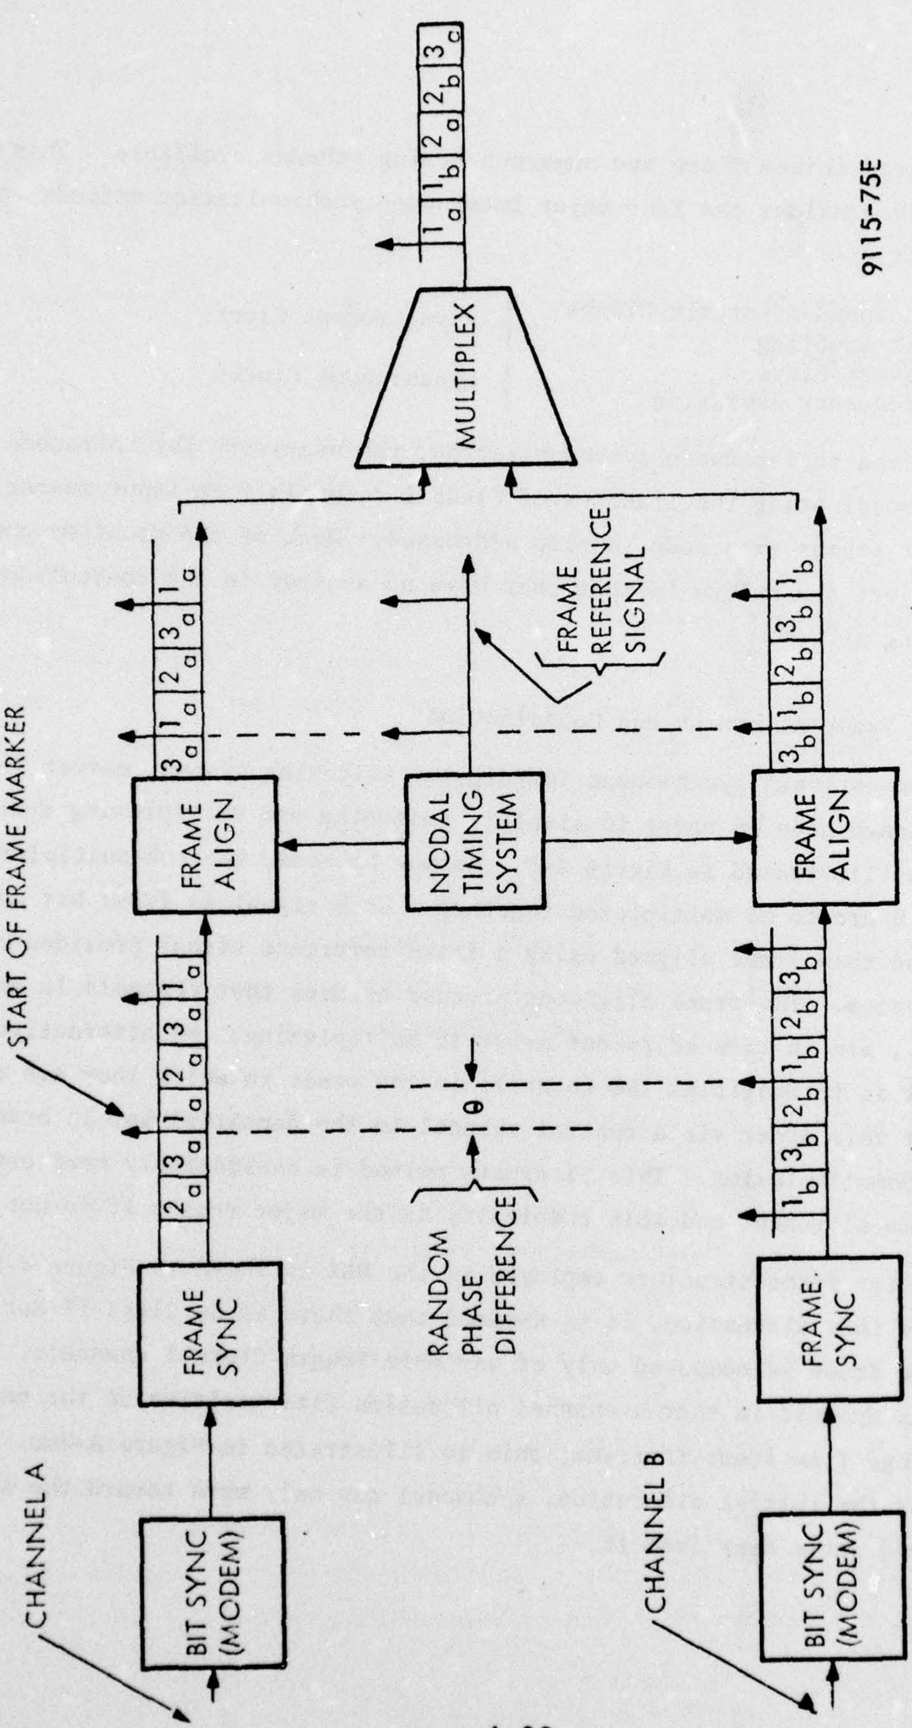

Figure 4-8. Conventional TDM System

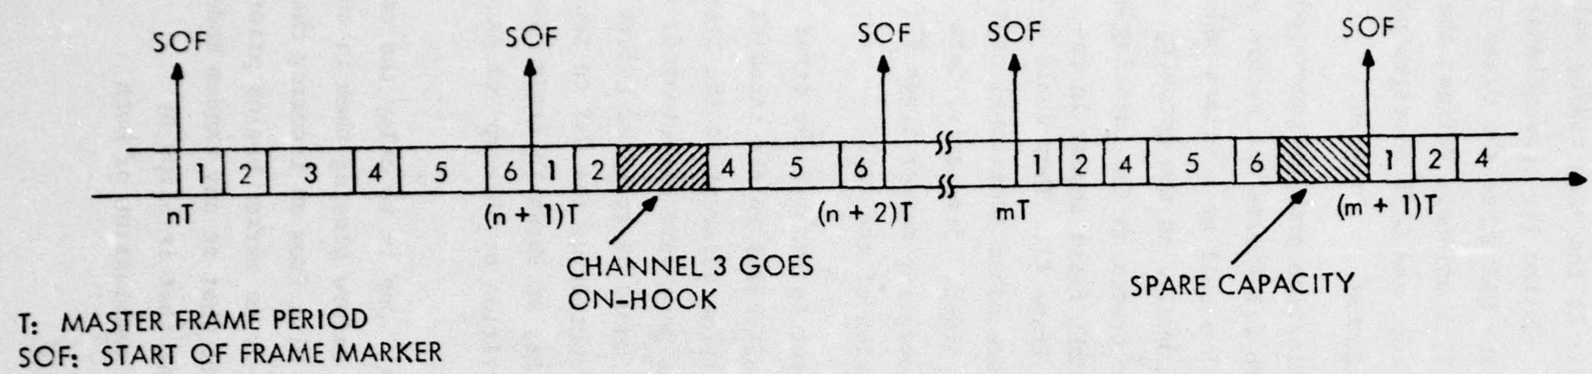

Figure 4-9a. DAX Master Frame Structure

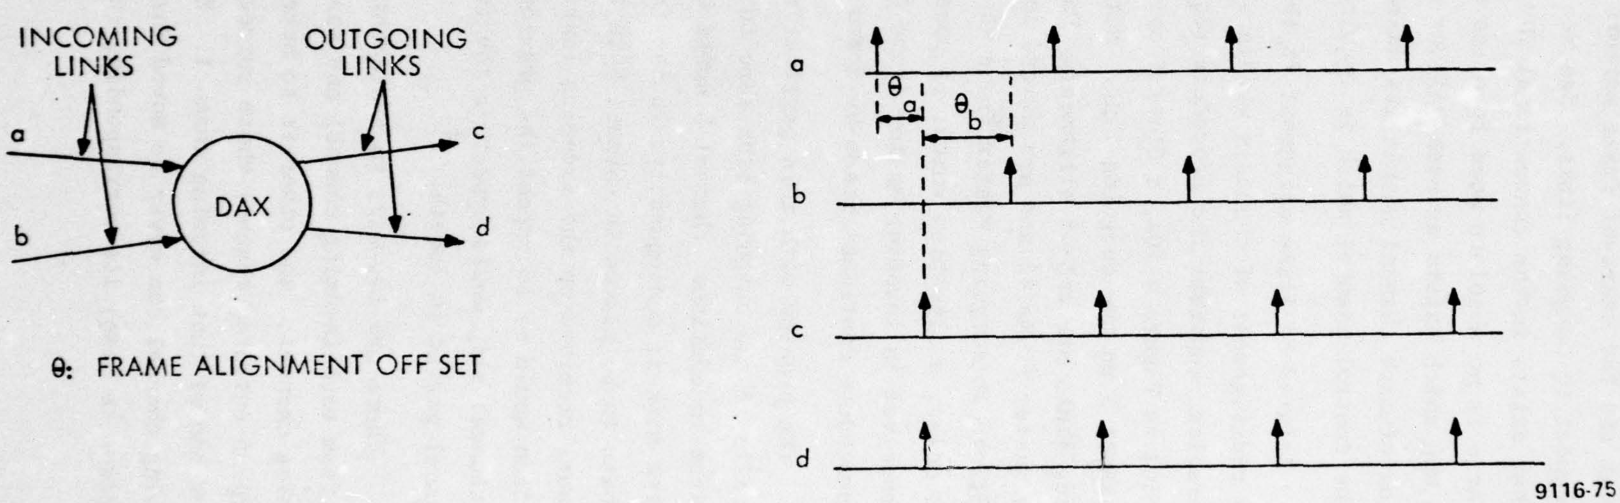

Figure 4-9b. Master Frame Alignment at DAX

At any DAX, the start-of-frame markers in all outgoing links are in time alignment and the start-of-frame markers in all incoming links are randomly aligned with respect to outgoing links. See See Figure 4-9b. This is the same timing relationship that exists in the conventional TDM described above. Unlike the conventional TDM, there is no absolute need to align master frames in the DAX since the Class I region map which exists at each DAX for each terminating link uniquely defines the location of each channel within its master frame. In effect, the CCIS messages provide the control channel which is required when frame alignment is not used.

Related to frame alignment is the timing or coordination problem generated by the reassignment of channels within a master frame. To illustrate the nature of this problem, consider the procedure required to establish a call at a tandem node. Referring to Figure 4-10a, a Class I call is allocated channel 6 on the incoming link and channel 5 on the outgoing link. Note that the master frames on the incoming and outgoing links are out of alignment. Data bits for the call first appear in incoming master frame #1 and are placed in outgoing master frame #1. They could have been placed in outgoing master frame #2 but this would have added additional cross-office delay, a situation which, if possible should be avoided. Similarly, data bits received in incoming master frame #2 are placed in outgoing master frame #2; this procedure continues frame-by-frame until the conclusion of the call.

The problem with this particular channel assignment is that if the first four calls in the outgoing link were to terminate while all calls in the incoming link were to continue, channel 5 would be moved to a position adjacent to the start-of-frame marker; consequently, in the frame in which this move were effectuated, the call bits to be placed in channel 5 in the outgoing link would be required before they were received on the incoming link. The probable reaction of the DAX to this situation would be to repeat the previous frames call bits, as shown in Figure 4-10b. This channel slip would appear to the customer as a repetition or 'hiccup' of data, a channel period in length.

There are two ways to avoid this 'hiccup' effect. One is to delay the call bits from each incoming channel one master frame period before placing them in an outgoing channel; the other is to never place the call bits from an incoming channel in an outgoing channel whose corresponding start-of-frame marker begins prior to the end of that incoming channel. Both methods assure that at any tandem mode an outgoing channel can never be moved forward to the extent that it slips by (precedes in time) its corresponding input channel. An examination of both

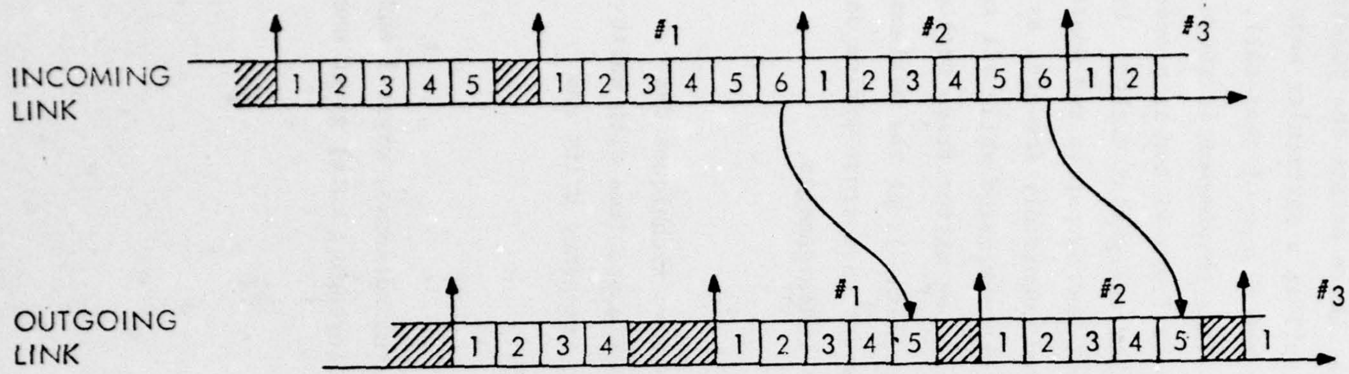

Figure 4-10a. Channel Allocation at a Tandem Mode

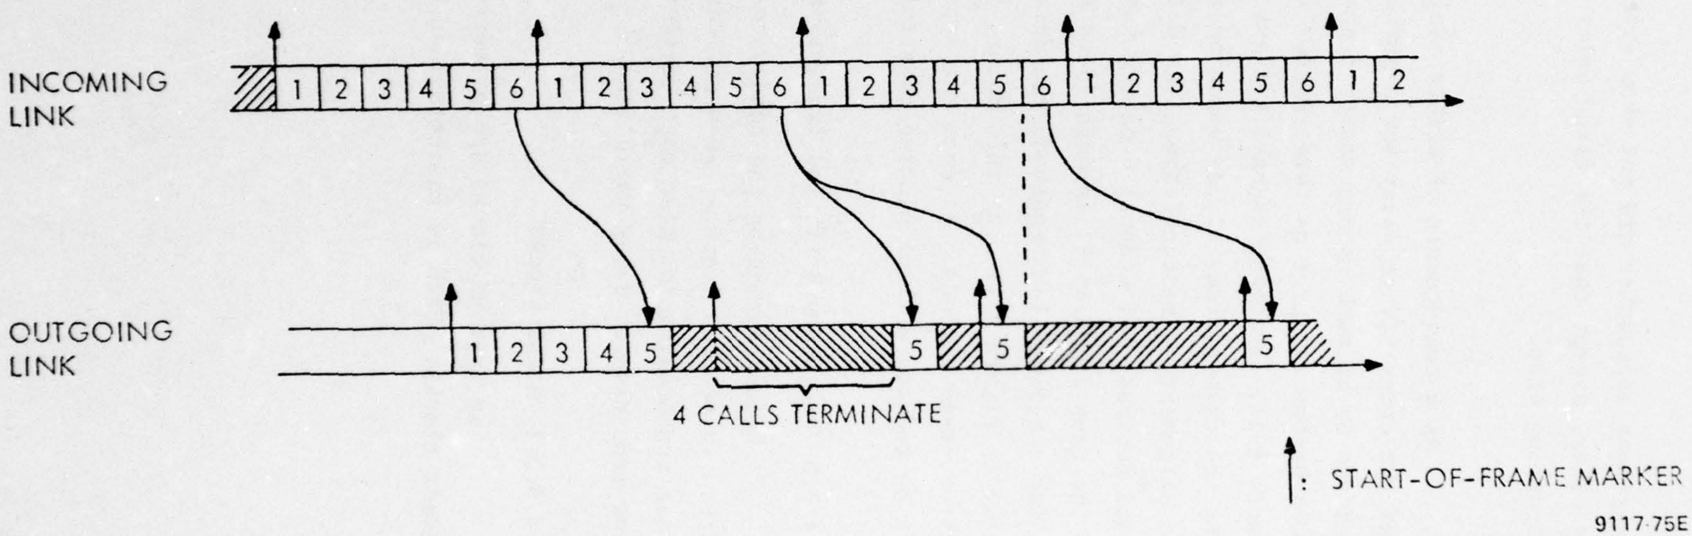

Figure 4-10b. Example of a Channel Slip

assignment approaches did not show either to possess a significant advantage over the other, except that the fixed delay approach appears to require somewhat less processing time.

As a consequence of either assignment approach, there exists the possibility that the cross-office delay for a particular Class I call at a particular node could increase by as much as one master frame period during the course of that call, and could become as large as three master frame periods. This phenomenon is shown in Figure 4-11. Incoming channel 6 first appears in master frame #1 and is placed in outgoing master frame #2, as would be required by either assignment method. Due to the alignment of master frames, the delay between input and output is two master frame periods. If channel 6 on the incoming link moved completely forward, as would be the case if calls #1 through #5 on the incoming link terminated while all calls on the outgoing link remained, the delay would increase one master frame period, as shown in Figure 4-11. At this point, the cross-office delay is at the maximum, three master frame periods. It does not appear that this worst case cross-office delay can be reduced without placing constraints on channel reassignments.

4.4.3.3 Frequency Drift and Transmission Delay

Before examining the candidate bit synchronization techniques described in Section 4.4.3.4, we must consider the two basic problems with which a bit synchronization plan must content: nodal clock frequency drift and link transmission delay variations.

4.4.3.3.1 Nodal Clocks

Two type of clocks are primarily used in digital networks; crystal and atomic clocks. Each is characterized by a stability (or aging rate) given usually in

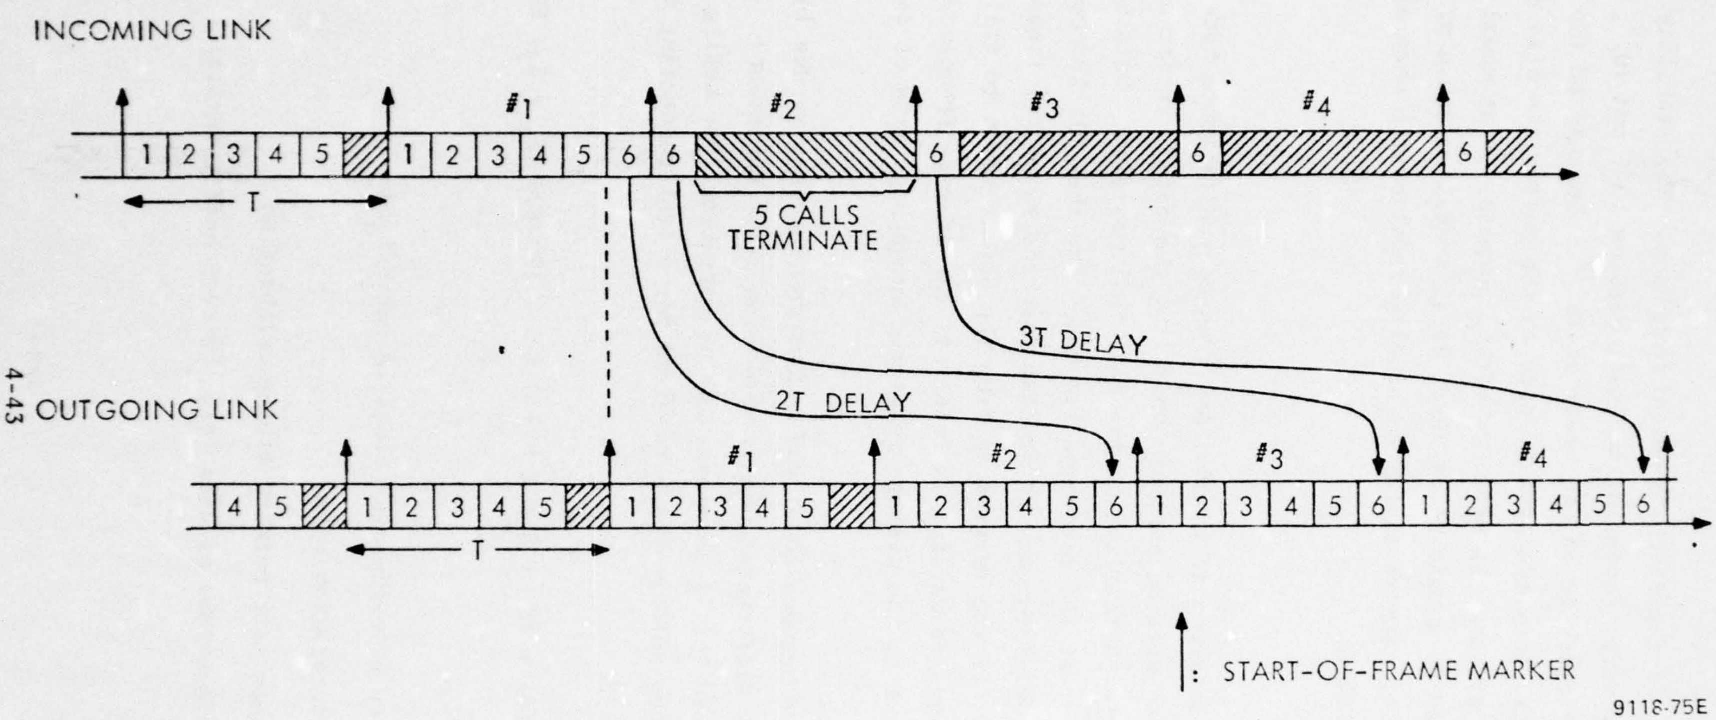

Figure 4-11. Worst Case Cross-Office Delay at a Tandem Mode

parts/f_o/day, where f_o is the clock's nominal frequency. The stability of a commercial crystal clock may vary, depending on cost, between 10^{-6} and 10^{-9}; the stability of a commercial atomic clock is on the order of 10^{-12}. Because of the greater stability of atomic clocks, they are usually used as the master clock(s) within a network. Crystal clocks are usually used for backup purposes or as nodal clocks (possibly phased locked to a master) at lower level nodes. Because all clocks exhibit frequency drift, network timing is usually recalibrated once or twice a year against a primary standard.

4.4.3.3.2 Buffers

When data is received at a node, any frequency difference between recovered timing and nodal timing will necessitate buffering the data. Buffer operation is as follows. The buffer is initialized to the half full position. Data is first clocked into the buffer at the recovered frequency and then clocked out at the nodal frequency. If the nodal frequency is greater than the recovered frequency, the buffer tends to empty; if the reverse is true the buffer tends to fill. Whenever a buffer overflows or underflows data is lost. If the average frequency of the nodal clock can be kept equal to the average recovered frequency, loss of data can be prevented.

The difference between nodal and recovered frequency is due to two sources, nodal clock frequency differences and transmission delay variations. For the present, it will be assumed that transmission delay variations are negligible. Subject to this assumption, the buffer size B required for a link connecting to independently timed nodes is given by

$$B = 2 \times 86{,}400 \times (S_1 + S_2) \times F \times C \times D = 172{,}800\ (S_1 + S_2)\ FCD$$

Where:

S_i: stability of each nodal clock in parts/f_o/day

F: link transmission rate

C: days since last network timing calibration

D: days before loss of data (time between buffer resets)

Implicit in the above formula is the worst case assumption that the nodal clocks drift in opposite directions. Figure 4-12 illustrates the relationship between buffer size and differential clock stability ($S_1 + S_2$) for the following case:

a. Network timing recalibrated every 180 days
b. Link transmission rate of 1.544×10^6 b/s
c. Buffers reset every 24 hours
d. Network to be recalibrated within one hour

4.4.3.3.3 Transmission Delay Variations

Transmission delay variations also cause timing problems; however, their effects manifest differently than do the effects of nodal clock variations. To see this consider two interconnected and independently timed nodes A and B. If the transmission delay between A and B increases (decreases) in both directions, then more (less) bits are stored in the link and the receiver buffers at each end of the link will empty (fill) simultaneously. On the other hand, if the nodal clock at A is faster (slower) than the nodal clock at B, then the buffer at A will empty (fill) while the buffer at B will fill (empty). It should be noted that by comparing buffer changes at both ends of a link, it should be possible to determine whether buffer changes are due to clock or transmission delay variations.

There are two common approaches for handling disturbances caused by transmission delay variations, inserting adjustable delays in the transmission path or allocating sufficient buffer space to absorb maximum delay variations. The latter method is simpler to implement but not as flexible as the former. Which of two methods is preferable depends on the magnitude and rate of change of the delay variations. Estimates of these quantities* for LOS radio and coaxial cable are as follows:

a. Coaxial Cable

  1. The dominant cause of delay variations is linear expansion cause by temperature change;
  2. Changes in transmission delay occur slowly;
  3. For a path length of 3000 miles and a transmission rate of 1.544×10^6 b/s, the change in the number of bits stored in a cable due to a 22°C temperature change is, approximately, 10 bits.

*Based on data in Reference (1).

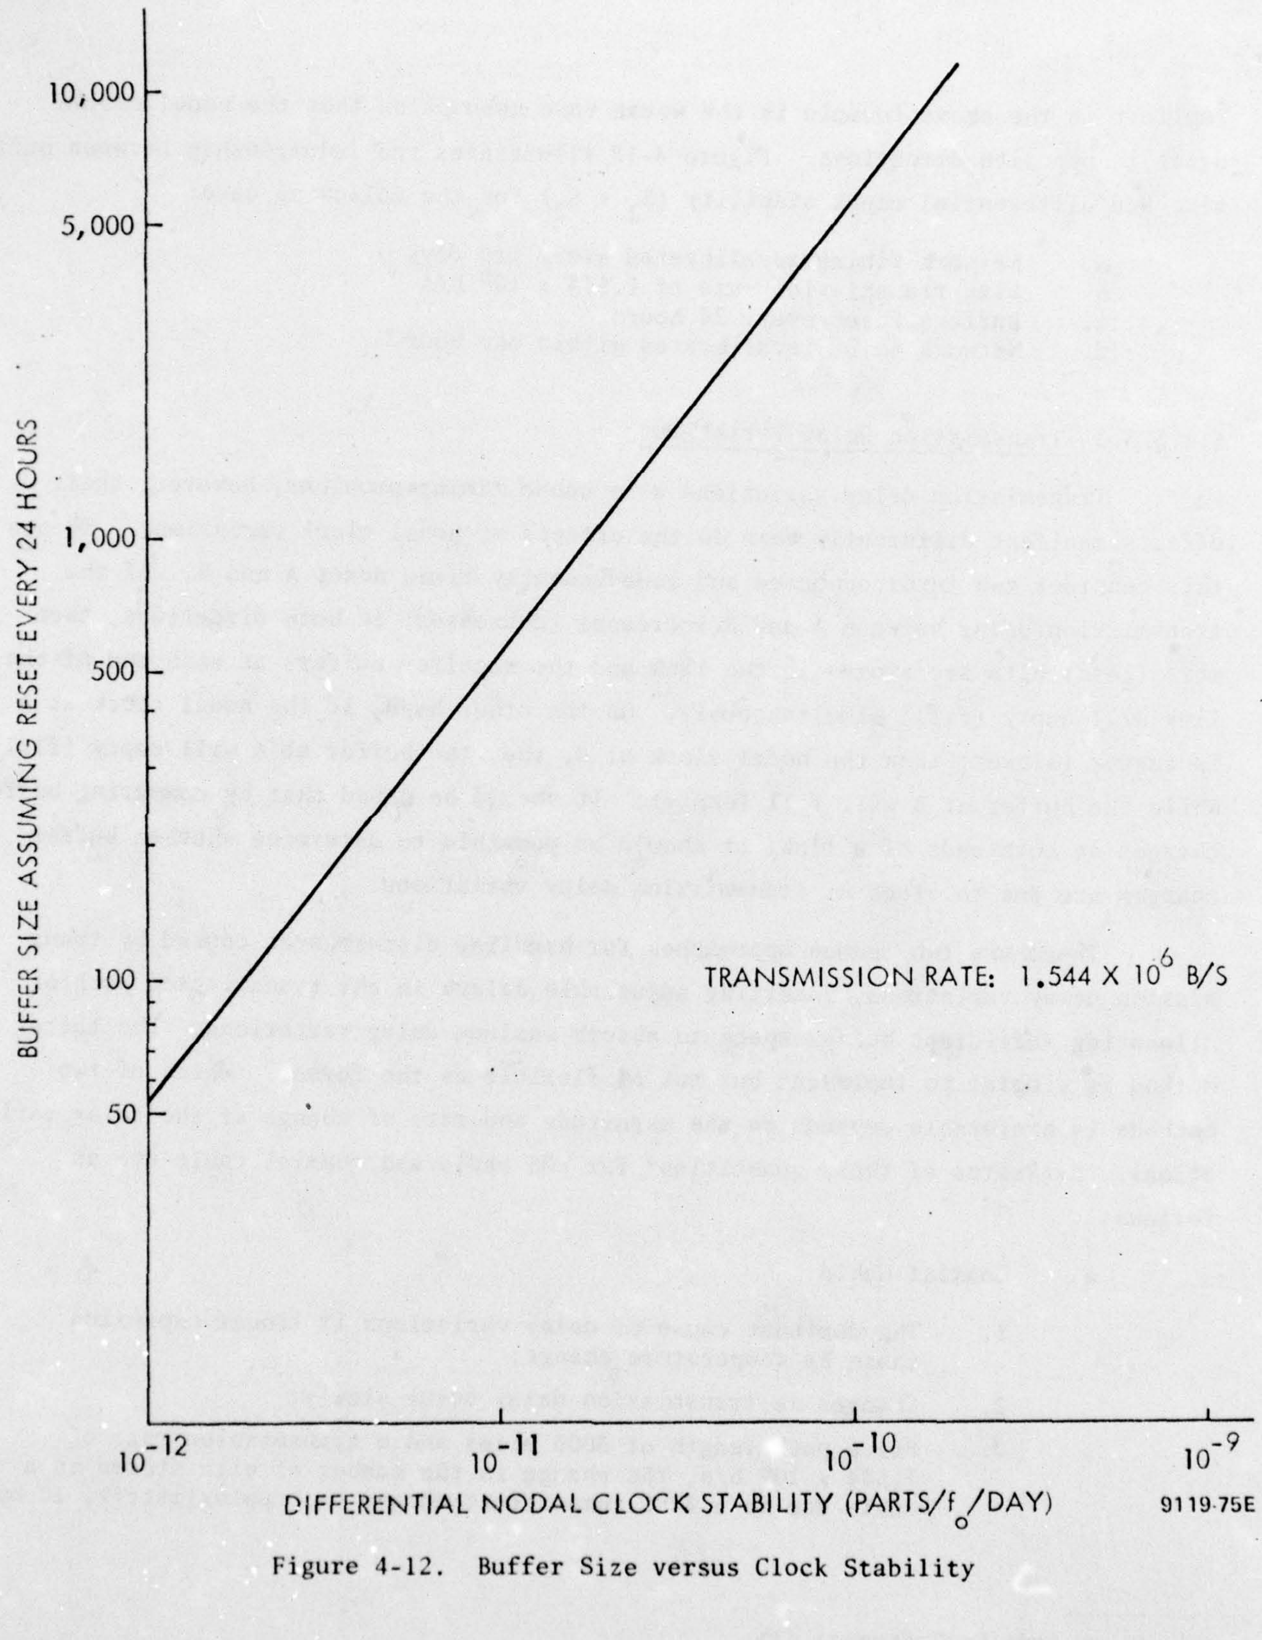

Figure 4-12. Buffer Size versus Clock Stability

b. Microwave Radio

1. Delay variations are due to changes in temperature, pressure, humidity and probably rain;
2. Rapid and long-term delay variations are possible;
3. Average variations in link bit storage are (1) daily, 0.49 bits; (2) monthly, 2.5 bits; (3) yearly, 7.4 bits.

It appears from these estimates that, at least for LOS radio and coaxial cable links, the less complicated and less costly buffer appraoch is more than adequate for absorbing delay variations.

With tropospheric transmission, short term delay variations are to be expected, but the magnitude of these variations are not significantly greater than those experienced in LOS radio links. Consequently, buffers, as opposed to variable delay elements should be used to absorb delay variations in these systems.

Delay variations in a satellite link during a period of 12 hours can be quite severe, due primarily to orbital eccentricity. Buffer requirements for such links on the order of 5000 to 6000 bits are possible. However, this buffering is provided by the transmission system and need not be considered a DAX requirement.

4.4.3.4 Techniques for Bit Synchronization

4.4.3.4.1 Independent Atomic Clocks

In this approach, each node in the DAX network is equipped with an atomic clock. During normal operation a crystal clock slaved to the atomic clock provides all nodal timing. If the crystal clock should fail, the atomic clock provides timing directly. If the atomic clock should fail, the crystal clock is slaved to one of the incoming links. If all links were to fail, network timing is not necessary since communications to other nodes is impossible; however, service to dedicated subscribers would continue in this mode using the now free running crystal clock as the nodal timing supply.

There are presently two types of atomic clocks commercially available, one based on cesium beam technology and the other on the rubidium gas cell technology. The cesium beam clock is a primary standard. It offers a stability of $\pm 1 \times 10^{-12}$ parts/f_o/day and at least that performance for longer averaging times. The rubidium standard offers approximately the same short term stability but exhibits finite long-term drift, typically $\pm 1 \times 10^{-10}$ parts/f_o/year. Based on these stabilities, DAX buffer requirements to compensate for rubidium or cesium clock frequency drift

can be computed using the buffer size formula given in Section 4.4.3.3.2. Assuming network recalibration every 180 days, buffer reset every 24 hours and a transmission rate of 1.544×10^6 b/s, the buffer lengths required† are given by

Rubidium: $172{,}800 \times (2 \times 10^{-12}) \times 1.544 \times 10^6 \times 180 \times 1 \sim 96$ bits

Cesium*: $172{,}800 \times (2 \times 10^{-12}) \times 1.544 \times 10^6 \times 1 \times 1 \sim 1$ bit

The chief advantages of independent atomic clocks over other internode synchronization schemes are

a. conceptually simple and a proven technique
b. good reliability
c. good survivability/low risk

The primary disadvantages include

a. moderately costly
b. periodic loss
c. atomic clocks are complicated devices and require trained personnel to maintain them

4.4.3.4.2 Bit Stuffing

In this method, all nodes possess independent, moderately stable crystal clocks. In order to prevent loss of data due to frequent buffer overflow/underflow, all data received at a node is retimed as follows. If an incoming bit stream is slower than the local clock, stuff bits are added to the bit stream to bring the two into synchronism. If the bit stream is faster, bits are removed or unstuffed and transmitted over a separate control channel. Extensive control signaling is required in order to unstuff and reassemble bit streams at each node.

Because digital switching occurs at each node in the network, all Class I channels would have to be stuffed and unstuffed individually at each tandem node as well as at the originating and terminating node. This requirement in conjunction with the dynamic master frame structure of the DAX makes bit stuffing probably unrealizable in terms of processing time and undesirable in terms of cost. This technique is therefore rejected and will not be given further consideration.

---

†Exclusive of transmission delay variation considerations.
*Assuming $\pm 1 \times 10^{-12}$ parts/f_o/180 days.

4.4.3.4.3 Master/Slave

In this synchronization plan, a particular node is designated master for the entire network and all other nodes slave their timing to this node. The master clock is usually a primary standard while all other timing supplies are moderately stable crystal clocks. For reliability, all clocks are made fully redundant. The choice for master node is usually based on geographical considerations. Network timing is disseminated via a timing tree whose links are a subset of the network links. The timing tree must be carefully maintained and fully redundant since loss of a single link can have catastrophic results.

The major objection to this scheme is that loss of the master clock severely cripples network timing. To illustrate this point, assume that the master fails and that the following optimistic assumptions apply:

a. All crystal clocks are at the identical frequency when the master fails and they free run at this last remembered frequency;

b. The drift rate is $\pm 1 \times 10^{-9}$ parts/f_o/hr;

c. Buffer capacity is 1000 bits.

Then using the buffer size equation in Section 4.4.3.3.2, the amount of time that passes before loss of bit integrity occurs on the average link is given by

$$1000 = 172800 \times (2 \times 10^{-9}) \times 1.544 \times 10^{6} \times D$$

$$D \sim 1.9 \text{ hours}$$

Because 1.9 hours is not considered a very large margin of safety and because of the vulnerability of such a scheme, the master/slave concept as described is not considered viable.

It is possible to reduce the risk involved in the master/slave method by modifying it in the following manner. The network is functionally divided into several smaller networks, each with its own master node. Synchronization in any subnetwork may be maintained as in the conventional master/slave concept or, as is usually the case, via a hierarchical timing tree. The latter approach requires that each node within a subnetwork be assigned a level and that the nodal timing at each node be obtained by slaving to a node one level above. At the top of the hierarchy is the master node; therefore, all timing within a subnetwork is directly or indirectly slaved to the master. Synchronization between subnetworks is assured by synchronizing the master nodes through atomic clocks or frequency averaging

This hybrid approach is considerably less risky than the straight master/slave scheme since it uses many masters clocks and since timing at any lower level node can theoretically be maintained by slaving to any incoming link.

4.4.3.4.4 Frequency Averaging

Frequency averaging requires that the frequency of a nodal clock be maintained at the average frequency or phase of all incoming links. This is accomplished as follows: For each incoming link at a node, a control signal is generated which is proportional to the buffer fill at that link. All control signals at the node are appropriately weighted and added together. This combined signal is then used to control the nodal clock frequency. For stability purposes, satellite links are permanently removed from the averaging process. Links which are out-of-service (out-of-frame, down, excessive error rates, etc.) would be automatically detected and temporarily removed from the averaging process.

This synchronizing scheme is attractive for many reasons

a. crystal clocks are adequate

b. no periodic buffer reset

c. relatively inexpensive

d. easily implemented using digital techniques and therefore easily maintainable

However, frequency averaging is not without its major drawbacks. Although theory predicts that the network will reach a single steady state frequency, frequency transients caused by natural (link failure, clock failure, start-up) or man made (jamming) disturbances ripple through the entire network. Studies have shown that transients caused by natural disturbances die out quickly and the system remains stable; however, jamming must be detected and counteracted in order to prevent catastrophic results.

4.4.3.4.5 Conclusions

Based on the foregoing discussions, it would appear that internode synchronization in the DAX network can be accomplished by any of the above schemes with the probable exception of bit stuffing. The choice of an 'optimum' scheme would depend on the criteria used. If vulnerability is the major concern, independent atomic clocks are the only choice. If this criterion is tempered by cost considerations,

the modified master/slave approach is a likely candidate. If cost and maintainability are the major considerations, frequency averaging appears very attractive. Network synchronization is considered further in Section 9.1.

4.4.3.5 Transmission System Flexibility

Up to this point, it has been assumed that all transmission links provide the same digital capacity. This Section will examine the impact on network synchronization of interconnecting nodes with links operating at various transmission rates, such as NATO group rates based on 256 kb/s or TRI-TAC group rates based on 576 kb/s.

With regard to bit integrity, there is no dependence on link transmission rate. This is especially true given the state-of-the-art in digital signal processing. As long as a single accurate and stable frequency source is available at each node, regardless of whether the source is an atomic clock or a crystal clock slaved to an incoming link, all required frequencies can be accurately synthesized with the same stability as the source.

The method of coordinating input and output channels at a node as described in Section 4.4.3.2 is independent of link transmission rate as long as the master frame period for all links is a constant. This is true because the only factor affecting the width of a Class I channel for virtual circuit switched operation is the period of the master frame. For example, assuming a master frame period of 10 ms, a 16 kb/s digital Class I connection would receive a 160 bit channel allocation in each link transversed by the connection regardless of the link's transmission rate.

When the master frame period is allowed to vary from link-to-link, channel coordination at a node becomes considerably more complicated. However, there appears to be no advantage to be gained by building this capability into the DAX, thus the implications of a non-uniform master frame period will not be considered.

4.5 COMSEC SYNCHRONIZATION

4.5.1 Problem

The SENET-DAX system must be capable of supporting encrypted digital signals from a wide variety of users. A major constraint imposed on the total switching and transmission system by typical modem comsec equipment is the need to establish and maintain end-to-end bit integrity.

4.5.2 Objectives

The objective of this task is to investigate the impact of providing bit integrity as a function of traffic class for a variety of user end-to-end configurations. Although the SENET-DAX may interface with other switching systems as illustrated in Figure 4-1 (Section 4.1) this study is limited to DAX local subscribers and to trunks between DAX switches. The detailed treatment of the necessary monitoring and control interfaces between modem comsec equipment and the DAX switching system is very complex and is not attempted at this time. Interest is focused primarily on those unique conceptual aspects of the DAX switching and multiplexing system that distinguish it from other digital switching systems such as the AN/TTC-39.

4.5.3 Analysis and Results

4.5.3.1 Cryptographic Synchronization Problem

The general nature of COMSEC equipment assumed to interface with the SENET-DAX system is characterized by the stream-cipher technique whereby a binary keystream is added bit-by-bit modulo two to each successive bit of plain text. The resulting encrypted digital sequence which appears to be random is transmitted to the destination where, in order to recover the plain text, it is necessary to add an identical binary keystream to the received sequence. If the keystream at the sending and receiving ends are not identical and/or are not synchronized the plain text cannot be recovered.

Encryption devices are provided for individual subscriber loops for handling digitized voice or data at a variety of rates and also for high rate trunk (or bulk) encryption. A secure voice or data call may be individually encrypted on a link-by-link basis or on an end-to-end basis. If the call traverses several trunks between DAX switches, some or all of the trunks may be bulk encrypted. From the point of view of an individual subscriber secure call, a single bit slip anywhere from the originating terminal to the destination can cause disruption of service. The transmission system design, the switch and multiplexing system design and the procedures and protocol for detecting and correcting synchronization errors must all function efficiently to minimize the amount of time lost to synchronization errors.

4.5.3.2 Trunk Encryption

Since trunk encryption impacts on all classes of traffic and on individually secure subscribers as well as non-secure subscribers, it is appropriate to consider this particular problem first. Trunk encryption of the master frame between DAX switching nodes may be necessary to provide traffic flow security. Such encryption will randomize the entire frame making it impossible to distinguish the Start-of-Frame Markers and the Class II packet envelopes in the event of interception of the encrypted trunk. Crypto synchronization is a prerequisite for obtaining and maintaining DAX frame synchronization. A similar state of affairs is common to other time division multiplexing systems. The DAX frame synchronization scheme differs from that used in the AN/TTC-39 and DGM families of digital multiplexing equipment in that it uses a multi-bit frame word to identify the start-of-frame rather than a single framing bit per frame. For maintaining frame sync and for signaling "loss-of-frame" or "request-for-frame sync," the DAX system uses a relatively short word on the order of 16 bits. For achieving master frame synchronization, the start-of-frame marker is a longer word, on the order of 48 bits. The use of multi-bit start-of-frame markers provides for rapid frame sync acquisition in the presence of random data on the remaining time division multiplexed channels. Thus the DAX concept does not require the setting to zero of all data channels during frame acquisition and the additional signaling necessary to coordinate the removal and replacement of data as in the above mentioned systems.

What is necessary of the DAX system is to provide frame sync subsystem performance under expected transmission bit error rate environments (including fading radio channels) that is adequate to provide low false alarm rates while also being able to rapidly and reliably detect true loss of frame sync. As in the AN/TTC-39 and DGM systems the procedures designed and implemented for the DAX frame sync management will differ depending upon whether or not a trunk encryption device is in the link and the nature of the transmission medium.

The loss of crypto sync on a trunk between switches will cause loss of traffic on all channels until crypto resync is achieved. The loss of sync must be recognized at the switch as a loss of master frame sync at the receiving end and the COMSEC equipment must be caused to enter a resync mode by an external command. The full duplex exchange of bits between COMSEC equipment to achieve re-synchronization must be reliable and typically requires several thousand bit intervals plus several round trip transmission delays. As soon as the COMSEC equipment has established crypto-sync, the traffic mode is continued using the long master frame marker. During the first few DAX master frames used for acquiring frame sync, Class II packet switched data and signaling may be transmitted over the trunk. Ideally the time delay from the loss of trunk crypto-sync until full duplex traffic is restored should be as short as possible compatible with the trunk bit rate, the transmission channel bit error rate characteristics, and the logical reliability required of the decision making algorithms. On some trunks, because apparent loss-of-frame sync may be due to a momentary high bit error rate, procedures can be provided to delay the crypto resync command by a fixed or selectable amount of time. If frame sync recovers during the waiting interval, the resync command is withdrawn.

4.5.3.3 Secure Voice Subscriber

As an example of a Class I user we consider the impact of crypto synchronization on a secure digital voice subscriber. For comparison we also consider a parallel non-secure digital voice subscriber. Assume for example that calls have been established and that voice communication is progressing normally. One of the more simple kinds of malfunction one might expect from a highly dynamic and

flexible multiplexing scheme such as DAX is that during a master frame when channels are being reassigned due to rapid changes in Class I traffic some individual user channel may occasionally lose or gain a bit. For most non-secure digital voice terminals, an occasional lost of extra bit would go unnoticed. If the bit integrity is disrupted on a secure channel, the subscriber loses crypto sync and hence voice communication. In this situation the switch itself may have no way of knowing that crypto-sync has been lost on a particular subscriber channel and simply continues to process the call. The user in this case must initiate the resync process. In the case of end-to-end encryption, as long as the switch maintains the connection, the COMSEC equipment at each user terminal can exchange the necessary on-line control signals to re-establish crypto-sync. Although end-to-end encryption involves fewer interfaces with COMSEC equipment and maintains continuity of message security, the procedures for key distribution throughout the entire switching network must be provided.

Link-by-link encryption can be used to connect subscribers having different area COMSEC keys and perhaps different terminal equipment. The DAX switching centers must be equipped with the necessary dedicated or pooled COMSEC equipment to decrypt and re-encrypt messages with the proper compatible crypto keys. The subscriber loops at each end switch may have crypto keys differing from each other and from those used on intermediate links between switches. The individual channel encryption is independent of any trunk or bulk encryption that may also be employed. When loss-of-crypto-sync occurs on a link-by-link encrypted call, the fault may lie in any segment and the problem of resynchronization is compounded. A crypto resync command initiated at a subscriber terminal must be recognized by the local DAX switch and a COMSEC sync procedure similar to that used in initially setting up the call must be exercised on each link even though all but one may already be synchronized.

On a voice call the subscriber becomes aware of the loss of crypto-sync immediately and can initiate action to re-establish the call. The user terminal for secure calls will have appropriate mode control buttons or switches which the user can manipulate. If resync procedures fail, the subscriber has the option to go On-Hook and later try to re-establish the call.

4.5.3.4 Secure Data Subscriber

The DAX Class I service may be used to support secure data links where the loss of crypto-sync may have to be recognized by a machine rather than a human. Many standard data formats provide means for detecting character errors and/or gross violations of message format. Loss of crypto sync would generally result in random bits not satisfying any prescribed format. If the data being transmitted over the switching network is of sufficient timeliness and importance, it may be necessary to provide redundant service over parallel or alternate routes. Alternatively all or parts of the data may have to be repeated with the aid of human intervention.

4.5.4 Conclusions

This has been on initial examination of the impact of crypto synchronization on the DAX system. The detailed specific requirements and procedures for interfacing with various existing COMSEC equipment is very complex. Certainly the synchronization and key distribution problems are of primary concern to users of secure links independent of the particular multiplexing scheme employed in the switch. Detailed examination of the interface requirements and operating procedures for specific families of COMSEC equipment must be integrated into the design of the DAX switch subscriber loop interfaces.

4.6 CCIS ERROR CONTROL

4.6.1 Problem

In order to control and coordinate the flow of traffic through the DAX network, it is necessary for DAXs to exchange control messages. These CCIS messages typically provide information relating to call originations and terminations, Class I channel reassignments, call preemptions, and synchronization, as examples. For reasons discussed in Section 2-2, CCIS messages will be in the form of data packets and will conform to ADCCP, a bit oriented data link control procedure.

ADCCP packets are inherently self-checking; that is, they utilize a 16 bit polynomial code for error detection. This task will examine the suitability of this form of error control with respect to the following DAX network parameters: type and average transmission rate of CCIS messages; noise environment; and retransmission scheme. The need to augment ADCCP error control with forward error correction (FEC) will also be discussed.

4.6.2 Objectives

There are three basic techniques available for error control in a data network:

a. Error detection and retransmission of the data found to contain errors (usually alluded to as Automatic Repeat Request, or ARQ);

b. FEC (errors remaining after the correction phase are passed on along with the good data);

c. FEC followed by error detection and retransmission (ARQ);

These three techniques along with possible variations offer varying degrees of reliability (probability of an undetected error) at the expense of throughput and complexity. The objective of this section is to examine these trade-offs with respect to error control techniques applicable to the DAX.

4.6.3 Analysis and Results

4.6.3.1 Background

Neglecting FEC for the moment, error control in the DAX network will be accomplished using the technique of continuous automatic repeat request When a packet is received at a DAX from a remote DAX, it will be tested for errors using the frame check sequence (polynomial code) contained within the packet. If no errors are detected, a CCIS control packet will be transmitted to the remote DAX specifying that it may erase from memory, subject to DAX accounting procedures per Section 3, the data packet in question. If an error(s) is detected, a CCIS control packet will be transmitted to the remote DAX requesting a retransmission of the erroneous packet. Retransmission of a packet will also occur if the packet "times out". That is, if a DAX does not receive an acknowledgement, positive or negative, in a specified time for a previously transmitted packet, the packet will be retransmitted. This "time out" applies to both control and data packets.

Because all DAX packets are sequence numbered, it will only be necessary to retransmit erroneous packets. Thus this particular ARQ implementation is superior to both stop and wait ARQ and complete retransmission/continuous ARQ* in that it provides greater throughput. For elaboration of this point see Section 10.3.2.

FEC with no ARQ does not provide sufficient reliability to be considered a viable alternative for CCIS error control; however, when used in conjunction with ARQ (henceforth denoted FEC/ARQ) it does merit consideration. In this latter context, each packet received at a DAX would first be FEC decoded using hardware and then error checked using either hardware or software. FEC is employed with the intention that for most packets it would correct all errors thus obviating the need for retransmission and, in turn, increasing throughput. It should be pointed out that this hybrid error control procedure is not necessarily superior to straight ARQ, as will be seen shortly.

---

* Complete retransmission/continuous ARQ requires that the system backup and retransmit the erroneous packet and all packets subsequent to that packet.

4.6.3.2 Noise Model

Error environment plays an important role in determining the error control to be employed on a channel. To illustrate this point assume the following applies to a typical inter-DAX link:

a. Bit errors are caused by a binary symmetric channel (BSC);
b. All errors are independent (random);
c. The BER is approximately 10^{-2}.

Then if the average CCIS control packet is 150 bits long, the average number of bit errors per packet is $150 \times 10^{-2} = 1.5$. Thus, this error environment appears well suited to FEC/ARQ since an average of 1.5 errors per 150 bits is well within the capability of most FEC codes.

On the other hand, if the bit errors on the channel occur in bursts and if errors within these bursts exhibit an average BER of 10^{-1}, then the suitability of FEC is not so clear. In this case, the average number of bit errors occurring in a packet caught in a burst is $150 \times 10^{-1} = 15$. Although there are many FEC codes which can handle 15 errors randomly distributed among 150 bits, there are likely to be many cases where the number of errors would considerably exceed 15 or where the errors would be bunched occurring in a span of say 25 to 50 bits. Although FEC codes can be designed to fit a particular burst noise model, it is rare that a real channel obliges by providing errors which fit the assumed noise model with any consistency.

Obviously a valid noise model is required in order to develop an adequate error control system. Unfortunately, error statistics applicable to wideband links of the size and multiplicity required to interconnect DAX's are not available. To circumvent this problem, the noise model employed in the AN/TTC-39 will be used. Since the AN/TTC-39 is a tactical switch probably comparable in capacity to a DAX, the DAX noise environment although possibly no better than the AN/TTC-39's, should certainly be no worse. The assumed noise model is already given in Table 4-3.

4.6.3.3 CCIS Packet Statistics

As described previously, CCIS messages are required for the efficient control of traffic. These messages will be transmitted in-band in the form of packets. They will be carried in the Class II region of the master frame, although during abnormal conditions they may be transmitted in the Class I region. A detailed description of CCIS message structure and procedures is provided in Sections 2.4, 3.1, and 3.2. CCIS packet statistics are derived in Section 10.2; the statistics pertinent to the problem at hand are as follows:

a. CCIS packet size varies from, approximately, 100 to 230 bits. The weighted average packet size is 150 bits;

b. The worst case equivalent BH channel required to handle CCIS packets is approximately 5500 b/s.

Based on these statistics and the assumption that the average inter-DAX link provides a transmission rate of 1.544×10^6 b/s, the following results are obtained:

a. An average packet period of (5500 b/s)/(150 b/packet) = 36.7 packet/s or one packet every 27.3 ms.

b. An average packet duration of (150 b/packet) x $(1/1.544 \times 10^6$ b/s) $\sim$ 0.1 ms.

It is interesting to contrast the burst noise model assumed with the equivalent CCIS channel described above. This is done in Figure 4-13. Note that the duration of a packet is short compared to the duration of a noise burst, regardless of the frequency of the noise burst.

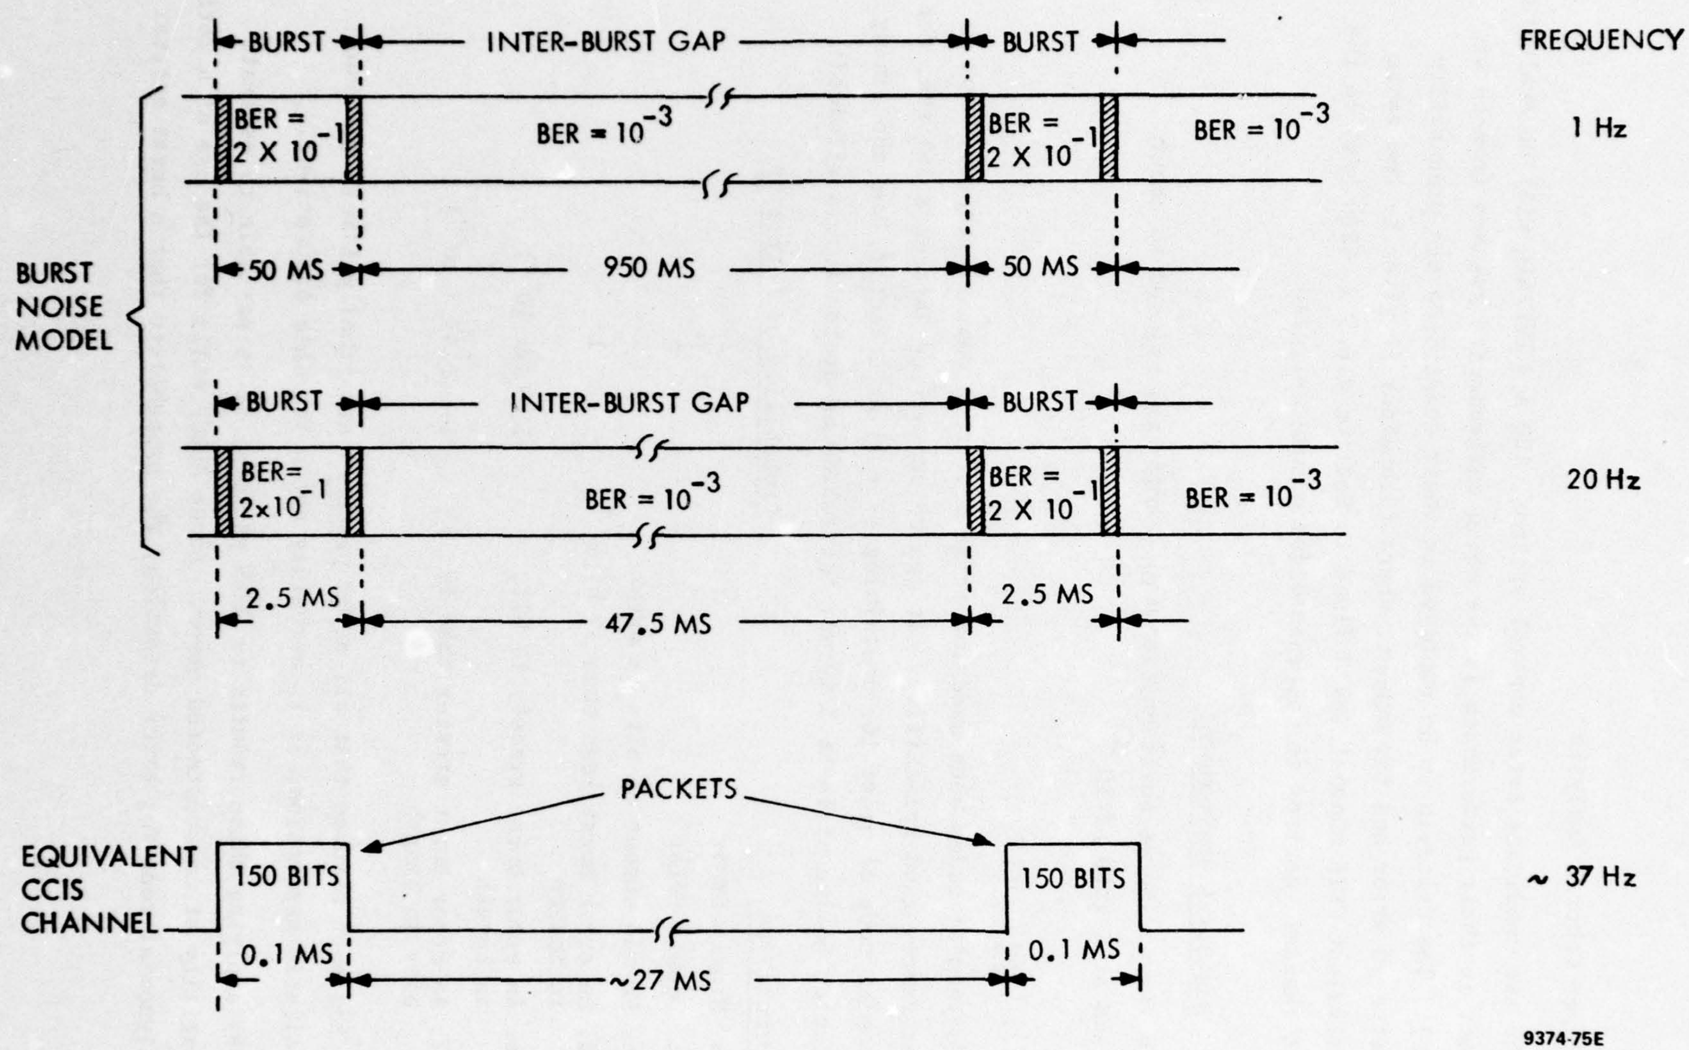

Figure 4-13. CCIS Equivalent Channel Contrasted With Burst Noise Model

4.6.3.4 Error Control Analysis

The two candidate error control systems, ARQ and FEC/ARQ will be evaluated with respect to their performance in the noise environments assumed (non-burst and bursts). The criteria to be employed in their comparison are probability of an undetected error and throughput, where throughput is given by the ratio of the equivalent CCIS channel (as defined in Section 4.6.3.3, 5500 b/s) to the equivalent channel required to compensate for channel errors.

4.6.3.4.1 Non-Burst Environment

In the non-burst noise environment, errors are assumed to occur randomly and at a rate of 10^{-3}.

4.6.3.4.1.1 ARQ

Polynomial codes when used for error detection can, if properly chosen, provide extremely good reliability. It can be shown (see Reference 36) that for any polynomial code of order 16, containing $(x + 1)$ as a factor, and containing another factor having at least 3 terms, the following protection is afforded:

Error Type	Probability of Detection
a. Single Error	1
b. Double Error	1
c. An odd number of bits in error	1
d. An error burst less than 17 bits in length	1
e. An error burst exactly 17 bits in length	$1- (3 \times 10^{-5})$
f. An error burst greater than 17 bits in length	$1- (1.53 \times 10^{-5})$

where it has been assumed that all error patterns are equally likely (although not a realistic assumption, it is certainly not a terrible assumption).

By applying these results to CCIS packets, it is possible to estimate the probability of an undetected error. Since ADDCP calls for the use of a 16th order polynomial code for error detection, the probability that a burst greater

than 17 bits in length goes undetected is 1.53×10^{-5} and the probability that a burst exactly 17 bits in length goes undetected is 3×10^{-5}. The probability of a burst greater than 17 bits in length in a received packet (150 bits) for the assumed random error environment is conservatively bounded by P (burst > 17 bits) $\leq \sum_{k=2,4...}^{150} \binom{150}{k} (.001)^k (.999)^{150-k}$

which follows from two facts: a burst must contain at least two errors; and all bursts with an odd number of errors are detected. This last equation may be evaluated as follows:

$$\sum_{k=2,4...}^{150} \binom{150}{K} (.001)^k (.999)^{150-k}$$

$$= \quad 1 - \sum_{k=0}^{1} \binom{150}{k} (.001)^k (.999)^{150-k} - \sum_{k=3,5...}^{149} \binom{150}{k} (.001)^k (.999)^{150-k}$$

$$= \quad 1 - ((.999)^{150} + 150 (.001)^1 (.999)^{149}) - (\binom{150}{3} (.001)^3 (.999)^{147} + ..)$$

$$= \quad 1 - (9.8987 \times 10^{-1}) - (4.76 \times 10^{-4} + ...)$$

$$\sim \quad 9.65 \times 10^{-3}$$

Similarly, the probability of a burst exactly 17 bits in length in a received packet is bounded by

$$133 \, (10^{-3})^2 \, (1-10^{-3})^{133} = 1.16 \times 10^{-4}$$

Therefore, the probability of an undetected error in a received CCIS packet is pessimistically upper bounded by

$$(1.16 \times 10^{-4}) \, (3 \times 10^{-5}) + (9.65 \times 10^{-3}) \, (1.53 \times 10^{-5}) = 1.51 \times 10^{-7}$$

To estimate the throughput obtained using straight ARQ, we observe that a packet is retransmitted only if it contains one or more errors. The probability of this event is given by

$$1 - (.999)^{150} = 1 - .861 = 0.139$$

Therefore, if n CCIS packets are transmitted from one DAX to another, the average number of retransmissions required after the first transmission of these n packets is 0.139 n. Neglecting those packets required to initiate the retransmissions, the total number of packets required to transmit without error the original n packets is given by

$$n\ (1 + .139 + (.139)^2 + (.139)^3 + \ldots) = n\ \left(\frac{1}{1-.139} \right) \sim 1.16\ n$$

which equates to a packet throughput of 1/1.16 = 0.861.

Another way to interpret this last results is in terms of the increase in the equivalent Busy-Hour CCIS channel required to offset channel errors. Using the worst case channel estimate 5500 b/s given earlier, the equivalent CCIS channel increases to 5500 x 1.16 = 6380 b/s.

4.6.3.4.1.2 <u>FEC/ARQ</u>

The impact of FEC in a random error environment is to effectively improve the BER. To illustrate this it will be necessary to assume an FEC code. The most suitable FEC codes to be used with DAX CCIS type messages are block codes as opposed to convolutional codes. Within the block code class, there are many codes from which to choose. It will be assumed that the 1/2 rate Golay code * will be used. This is a very powerful code and finds wide usage in this type of application.

If all CCIS packets are Golay encoded prior to transmission, then after reception and decoding the probability of a packet error is less than it would be without FEC coding. The probability of the second event is .139, as calculated in the Section above. The probability of a packet error when FEC is employed may be calculated as follows:

a. The probability of four or more errors in a 23 bit code word is

$$P_{ce}\ \sum_{k=4}^{23} \binom{23}{k} (.001)^k (.999)^{23-k} = 1 - \sum_{k=0}^{3} \binom{23}{k} (.001)^k (.999)^{23-k}$$

$$= 8.72 \times 10^{-9}$$

---

* The Golay code is a (23, 12, 3) code which means it can correct up to 3 errors in a 23 bit codeword derived from 12 information bits.

b. The probability of one or more of these erroneous codewords in a CCIS packet is approximately given by

$$1-(1-P_{ce})^{13} = 1-(1-8.72 \times 10^{-9})^{13} = 1.13 \times 10^{-7}$$

where it is assumed that the average CCIS packet encodes into 13 codewords.

As may be seen, FEC has decreased the probability of a packet error (one or more bit errors) from .139 to 1.13×10^{-7}. When FEC is used with ARQ it would be safe to say then that the probability of an undetected error is, for all practical purposes, zero. It should be recalled that for straight ARQ, the probability on an undetected error was conservatively estimated to be less than 1.51×10^{-7}.

If we make the simplifying assumption that with FEC the probability of a retransmission of a CCIS packet is zero, then the throughput realized by FEC/ARQ is still no better than 1/2, since FEC utilizes 1/2 rate encoding. This translates into an equivalent CCIS channel of 5500 x 2 = 11,000 b/s.

4.6.3.4.1.3 Discussion

For the random error environment just considered, there are two significant results:

1. Both ARQ and FEC/ARQ provide excellent reliability;
2. ARQ has the advantage over FEC/ARQ in throughput (0.862 to 0.5 for a BER of 10^{-3}). It can also be shown that as BER improves (BER < 10^{-3}), the ARQ advantage increases since FEC/ARQ throughput never exceeds 0.5. This is illustrated in Figure 4-14.

It would appear that ARQ without FEC is adequate in a random error environment. It should be noted though that if the BER increases past 10^{-3}, ARQ throughput decreases dramatically (See Figure 4-14). If such an occurrence is to be expected on a long-term basis, the decision to use straight ARQ would have to be re-evaluated.

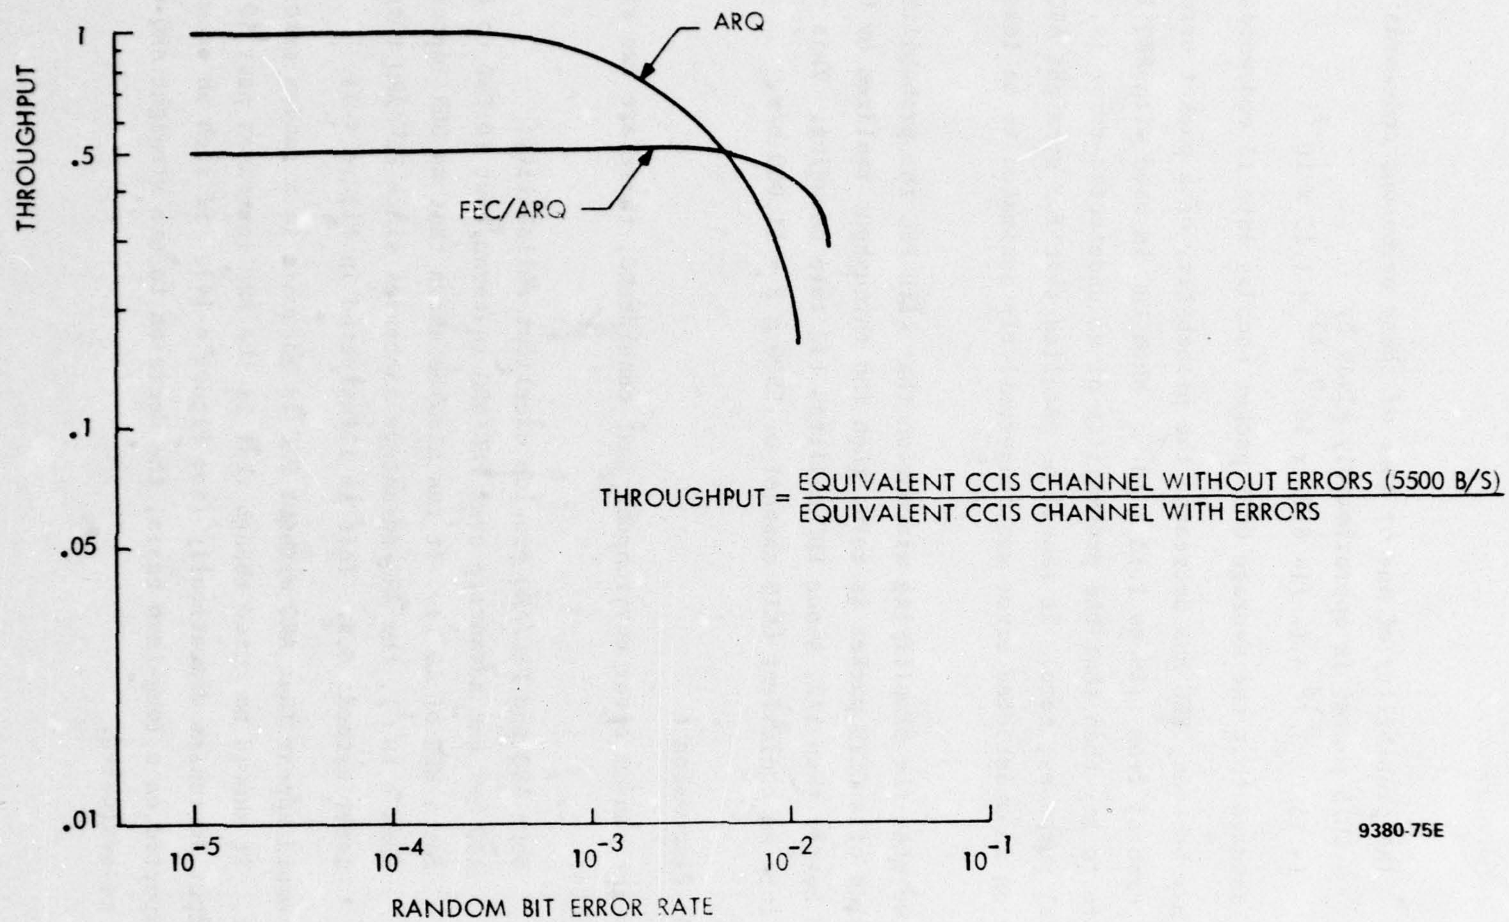

Figure 4-14. Throughput Versus Random Bit Error Rate

4.6.3.4.2 Burst Environment

In this environment, the errors occurring during a burst are assumed to occur randomly and at a rate of 2×10^{-1}. Between bursts the errors are also assumed to occur randomly but at a rate of 10^{-3}. The frequency of the bursts is uniformly distributed between 1 and 20 Hz and the burst duty cycle is 5%.

It can be shown that the conclusions reached for the non-burst environment also apply here. By referring to Figure 4-13 it may be seen that a CCIS packet because of its short duration as compared to a noise burst will in the great majority of cases either be hit by a noise burst (BER = 2×10^{-1}) or by an inter-burst gap (BER-10^{-3}) but not both. It also appears likely that most of the time a packet will be hit by an inter-burst gap rather than a burst, but this is of lesser importance. What is important though is the fact that when a packet is hit by a burst, the average number of errors expected is given by $150 \times 2 \times 10^{-1} = 30$. If straight ARQ is used, a packet caught in a burst will certainly have to be retransmitted. However, it is almost certain that a packet that is FEC-protected and caught in a burst will have to be retransmitted. Consequently, the only portion of the burst noise environment which impacts on the error control technique is the inter-burst gap which is, in fact, the same as the non-burst environment (BER = 10^{-3}). It follows then that the same conclusions apply; that is ARQ does not need to be augmented with FEX in this case.

SECTION 5

FUNCTIONAL ALLOCATIONS IN THE SENET-DAX SYSTEM

SECTION 5

FUNCTIONAL ALLOCATIONS IN THE SENET-DAX SYSTEM

5.1 SUBSCRIBER SERVICE FEATURES

5.1.1 Problem

Service features, as applied to telecommunication switching systems from today's available technology, are many and varied. To a great extent the application and deployment of the switch indicates the type of features to be implemented. In a commercial environment, the type of service provided for a PABX would vary from that of a central office, and the features of a central office would not be the same available for a tandem or gateway switch. Similarly, the deployment of switches within the military would indicate different types of features at each hierarchical level, or if not in type, certainly in quantity of application of individual features.

In designing a telecommunication switching system/network, it is necessary to take account of the constraints imposed by user requirements. The constraints imposed by the subscribers are determined by what ends the subscribers of the telecommunication system pursue. Some want to transmit an urgent message rapidly irrespective of network traffic loads or have some processing of messages carried out in the switching/transit centers. Other subscribers want to transmit large volumes of data without any code restriction. For others still, response time is all-important. Subscriber requirements can be summarized as follows:

a. Short transmission time, systematically or occasionally
b. Simplex, conversational, or inquiry/response-type communications
c. Transmission of large volumes of data
d. Temporary storage of data received outside office hours and automatic retransmission at office opening time
e. Automatic transfer of incoming traffic to an alternative terminal
f. Possibility of carrying out some simple processing in the switching centers
g. Low error rate
h. Security of transmission

i. Mode/code/speed/format conversions

j. Data operation facilities

The resulting constraints are numerous, varied, and often not easily reconcilable. This is definitely one of the reasons for the proliferation of custom-designed networks for the military and commercial world.

Two general classifications of service features can be delineated. These are subscriber-related features and system control/management features. This section is devoted to an assessment of the applicability and impact of subscriber-related service features for the DAX as an access exchange. The impact of subscriber-related service features for the DAX as a network nodal element in an integrated telecommunications network is discussed in Section 5.2. Assessments of system control/management features appear in Section 5.3

5.1.2 Objective

Many subscriber service features are available today in both circuit and message strategic and tactical switching systems. Several features useful for both classes of DAX subscribers have been identified from current switching systems (AN/TTC-38, AN/TTC-39, AUTOVON, AUTODIN, ARPA network), as well as those desirable with COMSEC operations. There are identities and similarities between the service features useful for both classes of subscribers, even though some features applicable to one class have no equivalence in the other. In some cases redefinition of the feature would make it serve both classes.

The objective of this section is to assess the desirability, feasibility, and impact of an optimized, cost-effective DAX being designed to offer particular services to data and voice-type subscribers of various classes.

5.1.3 Progress

Subscriber service features for each of the three main categories of switch services (circuit switch, packet switch, and store-and-forward) are assessed in terms of their impact on the DAX system architecture for both Class I and Class II subscribers. These features and their voice/data equivalents are listed in Table 5-1.

TABLE 5-1. SENET-DAX SUBSCRIBER SERVICE FEATURES

CLASS I	CLASS II
PRECEDENCE & PRECMPTION	PRECEDENCE & PREEMPTION
CONFERENCING	MULTIPLE ADDRESSING
INTERCEPT	INTERCEPT
DATA SERVICE	UNIVERSAL DATA SUBSCRIBER INTERFACE
SECURE MODE CONVERSION	MODE/CODE/SPEED/FORMAT CONVERSION MESSAGE/COMMUNICATIONS LINE SECURITY
COMPRESSED/ABBREVIATED DIALING	(NO EQUIVALENT)
CALL TRANSFER/FORWARDING	(NO EQUIVALENT)
AUTOMATIC GROUP HUNTING	(NO EQUIVALENT)
ATTENDANT/OPERATOR FUNCTIONS	(NO EQUIVALENT)
DIRECT ACCESS SERVICE	(NO EQUIVALENT)
(NO EQUIVALENT)	MESSAGE ACCOUNTABILITY
(NO EQUIVALENT)	ARCHIVAL STORAGE
(NO EQUIVALENT)	RETRIEVAL & TRACE

9944-75E

5.1.3.1 Class I (circuit-Switch) Service Features

There are approximately 10 basic subscriber service features available in varying degrees in strategic/tactical circuit switches today [23] and [60]. These service features include: precedence and preemption, conference capabilities, call transfer and call forwarding, compressed and abbreviated dialing, automatic line/trunk grouping and hunting, attendant services, direct access service, intercept, data service, and secure mode conversion.

5.1.3.1.1 Precedence and Pre-emption - DAX subscribers given this feature will be capable of establishing local, inter-switch, and extra-switch calls (with precedence compatible switches) using up to a 5-level precedence and pre-emption service. When a call (any type) is established, all connections are maintained at the precedence level assigned by the originator of the call, regardless of the precedence level authorized to other participants in the connection.

The five levels, in order of descending precedence, are FLASH OVERRIDE (FO), FLASH (F), IMMEDIATE (I), PRIORITY (P), and ROUTINE (R). Calls to switches employing no precedence, or non-compatible precedence levels, should be afforded the desired precedence level within the local switch and have the selected precedence translated to a level compatible with the switch serving the called terminal.

The operational concept of the precedence and pre-emption service should approach the following:

a. A call shall not be allowed to pre-empt a line, trunk time slot, or DAX function handling another call of equal or higher level precedence.

b. For a precedence call to a lower-precedence in-progress connection such as subscriber-to-subscriber, conference, trunk time slot, or subscriber participating in a conference, the existing connection shall be released and the called party shall be connected to the higher precedence originator.

c. Both subscribers involved in lower-precedence pre-empted calls shall be informed of the pre-emption by a "pre-emption tone", (or a pre-empt indication to the data subscriber) after the connection is broken. The called subscriber shall then be immediately connected to the

higher precedence caller, while the third party is returned busy tone for a specified time limit, after which his line is locked-out until he returns to the on-hook condition.

If the calling party requests precedence for a call, the routing control function can check his classmark for the maximum precedence level assigned to his line to make sure that it is not exceeded.

5.1.3.1.2 Conferencing - Four types of conference service are normally employed: progressive, preprogrammed, meet-me, and broadcast. Conference privileges can be assigned to subscribers through a classmark service, which is readily alterable. Various numbers of simultaneous conferences (of any type), of a fixed size, can be specified for a given DAX configuration. These numbers are generally linearly proportional to the switch size. When a conference consists of more than this fixed size, the number of independent, simultaneous conferences may be reduced accordingly, or several conference bridges can be interconnected.

Appropriate signalling from the subscriber is used to request a conference connection. If the calling party is properly classmarked, the processor connects the calling party to a conference bridge. Special numbers can be keyed to indicate a preprogrammed, meet-me, or broadcast conference. If a preprogrammed conference number is detected by the signal-receiving function, the called numbers are retrieved from the processor's memory. For a progressive conference, the calling party dials the number of the first conferee. In either case, the call processing functions proceed in their normal manner to ring the called party and connect him to a conference bridge. Broadcast conferees are also processed normally except that only the receiving pair of all called conferee lines is connected to the conference bridge. For a meet-me conference, the requestor is connected to a designated bridge.

5.1.3.1.3 Call Transfer and Call Forwarding - The operational concept for call transfer is as follows:

a. Any one of the assigned subscribers on a DAX is capable of having calls to his directory number transferred to another directory number.

b. The subscriber or an attendant is capable of initiating this service from his telephone instrument by first dialing a special access code within the constraints of the numbering plan of the DAX and then

dialing the directory number to which he desires calls to be transferred. This telephone instrument may still provide the capability to place outward calls while in the call transfer condition.

c. This feature is negated by the subscriber dialing the special access code and then dialing his own number. The negation is possible only from the subscriber's telephone.

For call transfer, the signal-receiving control function of the DAX can recognize the access code which is dialed by the party who wishes to have his calls transferred, then receive the numbers to which calls are to be transferred. A subroutine can check the calling line's classmark to be sure that he has transfer privileges, if these are restricted to selected subscribers.

5.1.3.1.4 Compressed and Abbreviated Dialing - Compressed and abbreviated dialing subscriber features are normally as described in the following paragraphs.

With abbreviated dialing, the switch completes local calls when the subscriber dials a portion of the full directory number, usually only the last few digits. Preemption may be effected by dialing the appropriate prefix digit(s).

Compressed dialing allows subscribers so designated to dial a special 2-digit number, plus a prefix or an "end-of-dial", to reach another subscriber. The DAX can translate the 2-digit code and route the call to the called party. If the calling party desires preemption service, he must dial the appropriate precedence digit first. The compressed codes may be applied to any directory number in the numbering plan as designated by the subscriber.

The signal-receiving control function of the DAX can recognize compressed and abbreviated dialing by examining the pattern of digits or by detecting the prefix or "end-of-dial". For compressed dialing, the program searches the compressed dialing table for the full directory number of the called party. For abbreviated dialing, no further action is required, because the program recognizes the digit pattern and infers the local area and office code during translation. In either case, translation, routing, and call completion proceed normally.

5.1.3.1.5 Secure Call Mode Conversion - The DAX may provide secure mode conversion and key change functions for properly classmarked loops and trunks employing incompatible voice security equipment by the use of crypto loop-arounds under control of the processor. Initial signaling needed by loops and trunks to establish the call may be transmitted and processed in the clear. Some voice security gear may signal in encrypted mode. A special dial code consistent with the numbering plan of the switch in conjunction with classmarking may be used in initiating a secure call to inform the switch of the need for secure call service and possible mode conversion or key change.

All secure calls requiring trunking may be routed over approved or appropriately encrypted trunks. The requirement for security can be conveyed from node to node via the CCIS channel.

5.1.3.1.6 Automatic Line/Trunk Grouping and Hunting - The DAX may provide automatic idle line hunting service for properly classmarked switch subscribers. Each group is usually capable of accepting from two to five subscriber lines. The features provided by this service are as follows:

a. The subscriber addresses in a hunting group need not be arranged in an orderly numerical sequence.

b. Hunting for an idle line in the group will occur whenever the called number in a group is busy.

c. In the event all lines in a group are busy, the calling party shall be returned a line busy tone.

d. When a pre-emptive call experiences an "all lines busy" condition within the called group and the calling precedence is higher than at least one of the busy lines, the lowest precedence busy line shall be pre-empted and the higher precedence caller connected to it.

A directory in the DAX processor's memory can store tables that list all members of each automatic line and trunk group.

5.1.3.1.7 Direct Access Service - The DAX should be able to afford designated subscribers with direct access service via the switched communication network. The switched direct access service can comply with the following requirements:

a. The direct access capability is designed for assignment to any of the subscriber terminals.

b. The scanning control function of the switch checks the classmarks of all lines originating a call to see if they have direct access service.

c. The switch, upon detecting an off-hook signal from an appropriately classed subscriber, immediately establishes a connection to a predesignated subscriber anywhere in the network.

d. The address and associated prefixes (e.g., precedence, route digit) are stored at the switching center.

e. The direct access subscriber is suitably protected against false seizure caused by misdirected control and routing signals and against non-seizure due to glare.

f. If an established direct access connection is disturbed for any reason, the connection can be reestablished by either one or both subscribers going "on-hook" momentarily, and then "off-hook" to initiate a new call. Alternatively "full-period" service will automatically re-establish a broken connection unless the release is detected at either party's terminal.

5.1.3.1.8 Attendant Recall/Operator Functions - The DAX may allow telephone subscribers to call or recall the local DAX attendant for assistance. The request is stored in the attendant's appropriate queue. The attendant extends either or both parties of a connection to any other termination to which they are normally allowed access. The attendant may extend calls or recalls at any level of precedence. Either connected party may release himself from the connection without releasing the connection between the attendant and the remaining party.

If the operator is busy at the time he is called, the calling party is placed in a queue that can be maintained by the software. When the operator releases, the first call of the highest precedence in the queue can be automatically connected by the software.

A special "release" subroutine can be used when the operator "hangs up". This will connect two parties on the bridge to each other. The two parties can now converse while the operator's bridge is available for other calls.

The operator's "hold" function can be implemented through use of the processor's software. Memory space can be provided to store those calls that are being held. The connections between the held parties and the operator's bridge are reserved. The processor can reconnect these calls in accordance with the operator controls.

Any other operator functions which may be required can be added by modular addition to the common control function.

5.1.3.1.9 Data Service - The DAX can provide real-time data switching for certain bulk data subscribers (See Section 2.2). This service is provided to properly classmarked lines of the switch who require large-volume transmission over long periods of time (e.g., sensor data). Data service can be provided for terminals operating with a variety of signaling and supervision plans.

The DAX processor must recognize when a Class I bulk data call is being processed so that it can select functions capable of handling the data rate of the calling party, compare sending and receiving data terminals for compatibility, and inform the called party that he will be receiving data instead of voice. If data service is requested, the program examines the calling line's classmarks to determine data terminal characteristics, and allocates time slot channels accordingly.

5.1.3.1.10 Intercept - An intercept feature provides for the return of a recorded announcement or connection to the operator when the number dialed does not exist, is unassigned, or is marked out of service.

Intercepted calls may be connected to the operator or they may be connected to a recorded announcement. If connected to the operator, they can be processed as described for calls to the operator after the intercept condition is recognized. If connected to a recording, they can be connected to a bus, similar to a tone bus, by the call completion program.

5.1.3.2 Class II Packet Switch and Store-and-Forward Service Features

There are approximately nine basic subscriber service features available to varying degrees in strategic/tactical packet/store-and-forward switches today.[(4),(22),(54)] These service features include: message accountability, message and communications line security, connection/interruption by precedence and pre-emption, multiple addressing, universal data subscriber interface capability, mode/code/speed/format conversion, archival storage, intercept, and retrieval and trace.

5.1.3.2.1 Message Accountability - Accountability in packet and store-and-forward switching concerns the establishment of very low probabilities that a packet or message, once acknowledged by a switch, will be lost or not transmitted, and to equally low probability that a legitimate packet or message will not be recognized on input, and therefore not serviced.

The DAX operational concept calls for messages to be broken up into packets of a fixed length, optimized for network throughput. These are the Class II data transmissions. Note that an exception exists for large quantities of bulk data, which will be treated as Class I data transmission (see Section 2.2).

The packets forwarded from node to node must be stored in the sending node until they have been acknowledged by a special control indicator. Within the packet envelope a numbering scheme and frame check sequence are provided. The numbering scheme gives an identification number so that the packet is uniquely registered and messages can be reassembled at their terminating node. The frame check sequence allows the receiving node to detect transmission errors.

In case a transmission error is detected, the incorrect packet is discarded. The sending node, after having received a negative acknowledgement link control indicator, or a time-out without having received a positive acknowledgement, can repeat the incorrect or any subsequent packet.

When a DAX has accepted and acknowledged a packet it is responsible for that packet until it in turn receives an acknowledgement from the next DAX on the transmission route, in which event it can discard the packet. Thus, tandem DAX's need only retain the packet long enough to determine that the packet has been received by the next switch. Full message accountability is maintained by the originating and terminating DAX's.

5.1.3.2.2 Message and Communications Line Security - Digital transmission facilities of the future DCS will be encrypted with equipment of the TENLEY type. Several possibilities exist for the protection of traffic from the data terminal equipments (DTE) to the DAX. If both the DAX and the DTE being serviced are in a secure area, or connected via approved circuits, then no COMSEC equipment is needed on the DTE loops. In addition, where COMSEC equipment is located at the DTE, it would be possible for the DAX, via its link protocol, to transfer encrypted traffic without decrypting it at the DAX. If netted crypto variables are utilized for local secure calls, it is also conceivable that the DAX may connect an appropriate device upon receipt of a distinctive signal. It may also prove practical to locate some TENLEY AKDC/LKG capabilities at the DAX.

5.1.3.2.3 Precedence and Pre-emption - As discussed in Section 3.2 it is our present intent to equate corresponding precedence levels of Class I and Class II traffic (shown below) in establishing the rules for pre-emptive traffic. The operational concept for these pre-emptive rules is discussed in Section 3.2.

EQUIVALENCE OF CLASS I/CLASS II PRECEDENCE

Class I (voice)	Class II (data)
--- (none)	W (Critic)
FO (Flash Override)	Z (ECP)
F (Flash)	F (Flash)
I (Immediate)	I (Immediate)
P (Priority)	P (Priority)
R (Routine)	R (Routine)

As the system architecture evolves, the intent is to minimize the complexity of these rules in order to minimize the impact on the DAX software and hardware as well. It is also intended to minimize the processing load (time burden) while providing a reasonable pre-emptive service.

5.1.3.2.4 Multiple Addressing - Multiple addressing with single input will have a significant impact on DAX storage and processing capability. A network handling multiply-addressed messages will have to work out some alternative to the source-to-destination flow control (i.e., not transmitting until receiving-storage is guaranteed at the destination) that works well for point-to-point traffic as in the ARPA network.

In such a network, one would neither want to repeat the message individually to each destination, nor (probably) to suspend transmitting it until all destinations had guaranteed storage. The likely alternative is to provide more buffer storage at nodes--significant high-speed storage for short, top-priority messages, and disk or tape--slower but cheaper--for longer messages that can tolerate more delay. A network of this type would thus combine packet switching with customary store-and-forward operation.

5.1.3.2.5 Universal Data Subscriber Interface - The DAX should be capable of interfacing with a family of universal data adapters. The TRI-TAC Data Adapter (DA) is representative of this family. The DAX should be capable of full-duplex operation with the DA at a variety of loop and data rates [59]. The DAX should be capable of operation with the family of DA's directly on a dedicated basis using the Dedicated Loop Encryption Device (DLED) or on a switched basis using a Digital Subscriber Voice Terminal (DSVT).

This interface would impact the transmission control hardware/software architecture of the DAX. Message information transfer between the DAX and the DA (and between DA's) would be controlled by means of a Data Adapter Control Block (DACB), which is used to establish the parameters of the ensuing data transmission. The processing of these parameters and the subsequent synchronization/control protocol must be considered in design of the DAX interfaces.

5.1.3.2.6 Mode/Code/Speed/Format Conversion - The DAX should be capable of interoperability with a wide variety of Class II data terminals/sources. Table 5-2 summarizes the various types of modes/codes/speeds/formats with which the DAX should be compatible. The extent of the mode/code/speed/format conversion required will certainly impact DAX architecture.

5.1.3.2.7 Archival Storage - It is suggested that regional DAX or AUTODIN type centers have the capability for archival storage of large amounts of bulk data for predetermined times. Through the use of a data communication protocol this archival bulk data may be accessed by any DAX in the network.

The impact of this feature is primarily the bulk memory storage required at the regional DAX centers. Also, access and retrieval data communication protocols must be handled by the DAX control function/processor.

TABLE 5-2. TYPES OF MODES/CODES/SPEEDS/FORMATS

MODES	CODES
• MODE I - CONTINUOUS/BLOCK-BY-BLOCK	• ASCII
• MODE II	• BAUDOT (ITA-2)
• MODE III - CONTINUOUS/BLOCK-BY-BLOCK	• CONT. RANDOM BIT STREAM
• MODE IV - RECEIVE/TRANSMIT	• 4 OUT OF 8 CODE (IBM)
• MODE V	• EBCDIC
• MODE VI	• FIELDATA
	• BINARY NONSTANDARD PARITY MAG TAPE (AUTODIN)
	• BINARY STANDARD MAG TYPE (AUTODIN)
	• DATA FORMAT MAG TERMINAL (AUTODIN)
	• FAX

SPEEDS (BPS)	FORMATS
• 8000N $N = 1,2,4$	• ACP-127
• 75×2^N $N = 0,1,2...7$	• JANAP-128, DATA
• 45.5	• ACP-127, MODIFIED
• 50	• JANAP-128, TELETYPEWRITER
• 2000	• NO FORMAT
• 4000	• SPECIAL FORMAT 1
• FUTURE HI-SPEED FAX	• SPECIAL FORMAT 2

5.1.3.2.8 Intercept - Among the more common message handling requirements in a strategic/tactical message switching environment is the capability to intercept temporarily all traffic for specific routing indicators. This necessity typically occurs due to the mobility of subscribers and the occasional difficulty in maintaining communication with stationary subscribers in tactical environments. When these situations are encountered, the normal message switch response, upon supervisory command, is to intercept any traffic that would otherwise be delivered to a temporarily inaccessible subscriber; to store this traffic temporarily; and when informed administratively that the subscriber can be reached again, to deliver to that subscriber any received traffic that has been stored.

The capability of executing this function would impact the DAX memory storage. One possibility is to have a set of regional intercept centers with enough bulk storage to satisfy the intercept function for all the DAX's in a given region. Intercept data communication protocols would then have to be developed to remotely store and access from these regional intercept centers.

5.1.3.2.9 Retrieval and Trace - Among the capabilities offered to subscribers of strategic/tactical message switches are the capabilities to trace or retrieve a message. Trace alludes to the capability to specify a message (by various identification criteria) and to receive in return a notification of the time of receipt or time of delivery of the message. Retrieval is a capability wherein a copy of a similarly identified message is retrieved from archival storage, and upon supervisory approval, is transmitted to the requestor.

It is suggested that these capabilities could be made available to the DAX in that the retrieval and/or trace messages would flow to the DAX which would automatically route them to an associated regional center (such as an AUTODIN center), and would similarly route the returned trace and/or retrieved message back to the subscriber, as directed by the switch supervisor at the regional center. Other than a delay measurable in seconds, there would be no apparent difference to the data subscriber in a retrieval or trace accomplished via a DAX, and a similar action taken as a dedicated or switched direct subscriber of a regional AUTODIN center.

5.1.3.3 Relation of Subscriber Service Features to Switch Size

In practice, the service features described have shown an increase in kind, magnitude, and sophistication as switches have grown in size. This is often because the "cost" of a feature (in complexity or dollars) is usually not directly proportional to the number of terminals serviced, and thus this cost can be amortized over more subscriber and trunk terminations. Today's developing technology leads us to believe that the "cost" of features will decrease such that more varied and sophisticated features will be practical in smaller switches. However, the relative scaling of features among switches of various sizes will probably continue.

For illustrative purposes, Table 5-3 lists the type and quantity of Class I service features according to switch size for three tactical military switches now in development and production by GTE Sylvania.

An initial assessment has been made of the impact of switch subscriber service features on DAX system architecture. With the possible exceptions of digital voice conversion, which remains to be explored further, it appears thatthese features can be provided within the SENET-DAX concept. Each feature contributes to software complexity (queue structure, list processing variations, etc.), to processor loading and consequent throughput, and to interswitch signaling time. Additional program, random-access, or bulk memory at each switch is often required, as well as the provision for digital control signaling between line and trunk terminations and the switch itself. The overall impact of each feature in SENET-DAX system software and hardware architecture must be evaluated in conjunction with the requirements of other service features, taking full advantage of commonalities among them.

Besides the general impacts described above, there are specific switch hardware impacts associated with many subscriber features. These include conference bridging facilities, operator/attendant positions, other common equipment for voice intercept announcements or signaling, and COMSEC transmission facilities. System tradeoffs must be made on the use of bulk or core storage, and optimal buffer utilization, for Class II multiple addressing and intercept. The provision for bulk archival storage leads to consideration of whether this can be provided at the originating (or terminating) DAX, or whether it is better provided at some regional center.

TABLE 5-3. SERVICE FEATURES VERSUS SWITCH SIZE

SWITCH DESIGNATION	SB-3614	AN/TTC-38	AN/TTC-39**
Number of Terminations	30-90	300-600	300-2400
1. Precedence and Pre-emption	2 levels Priority and Routine	5 levels	5 levels
2. Conferencing			
a. Progressive	2 at 10 parties	4-5 party conference bridges	10-5 party conference bridges per 750 lines; max. of 20 parties per conference
b. Preprogrammed	No	1-9 party conference bridge: 10 lists of conferees	20 list per 750 lines; max. of 20 conferees per list
c. Broadcast	No	No	Max. of 32 conferees
3. Call Forwarding	No	Yes, up to 25 subscribers	Yes, max. of 40 subscribers per 750 lines
4. Compressed Dialing	No	No	2 digit + end-of-dial; up to 100 subscribers per 750 lines may have 16 codes each
5. Abbreviated Dialing	Yes	Yes	Yes
6. Call Transfer	Yes, via attendant recall	Yes, via attendant recall	Yes, via attendant recall
7. Secure Call Mode Conversion	No	No*	Yes
8. Automatic Line/Trunk Grouping	2 groups of 4 trunks	Line: 30 groups with up to 5 lines per group. Trunks: 120 groups of any size	32 groups, 2 to 5 lines (terminations) per group
9. Attendant Services	Yes	Yes	Yes
10. Data Service	No	No	Yes, Real-time variety of interfaces
11. Intercept	To Attendant	To operator, information attendant, or error tone.	Recorded announcement

*The AN/TTC-38 is equipped to recognize requests for and connect subscribers to wideband trunks.

**AN/TTC-39 Features are as called out in Reference [60].

With regard to mode conversions, the possibility certainly exists for the DAX to provide translation, as well as key change functions where required, for connections between incompatible voice equipment. Appropriate control information would be conveyed via CCIS procedures and signaling. Obvious impacts exist in the areas of translation needs in software, as well as in the possible use of hardware conversion devices utilizing microprocessors or other integrated circuitry.

5.2 NETWORK IMPACT OF SUBSCRIBER FEATURES

5.2.1 Problem

As described in Section 5.1, current circuit switching, packet switching, and store-and-forward systems offer varied services to their respective types of subscribers. Some subscriber service features have an impact only on the local DAX switch equipment and processing, while others have a measurable impact on network performance and equipment that must be evaulated.

5.2.2 Objective

This section is devoted to an assessment of the impact of subscriber-related service features for the DAX as part of an integrated network. There are two main types of network impact which are evident. One of these is the impact on the switch itself, and the other is the impact on the network with the switches considered as nodes in the network configuration. These are somewhat interrelated; however, this section will primarily look at the DAX network, and only secondarily at the DAX itself.

5.2.3.1 Definition of Network Impact

By the term "network impact", we are referring to a measurable effect on the network, indicating a change in performance of some kind that can be shown to occur because of the use of a feature, or when network equipment must be provided (e.g., COMSEC transmission facilities) that would not normally be required. This effect, or measure, is a way of evaluating how a certain subscriber-related service feature enhances or restricts the performance of the network.

Specific measures of performance are usually related to the capacity, speed and "cost" of the network as affected by the individual services of interest. Measures thus far identified are:

a. Traffic Throughput: number of non-blocked voice calls and packets (excluding overhead traffic) existing simultaneously in the network

b. Network Grade of Service: fraction of offered traffic in the entire network that will be rejected before a connection is established, or delayed before data is transferred

c. Link Grade of Service: fraction of offered traffic between two switching nodes that will be rejected or delayed

d. Call Completion and Data Transfer Speeds: time needed to establish a voice call, from end of calling terminal signaling to beginning of notification to called terminal, or to transfer a data packet through the network

e. Network Control Complexity: this determines the "cost" of a network, as measured not only in design and implementation dollars, but in maintenance needs and modification possibilities and in the often intangible factor of subscriber acceptance

f. Special Equipment: This includes COMSEC transmission equipment for line/trunk security and data adapters for interface flexibility.

It should be noted that a comprehensive estimate of network performance cannot usually be made until subscriber density, available trunk facilities, and similar descriptive factors are considered in the analysis.

5.2.3.2 Impact of Class I Service Features

5.2.3.2.1 Precedence and Pre-emption - This service provides preferential treatment of a call by the network, whereby calls to busy stations will pre-empt a conversation existing at a lower precedence level. Multi-level precedence protection of calls and pre-emption of lower precedence calls when idle circuits are not available provides privileged subscribers a high call completion assurance, tending to decrease trunk use due to retries. However, the burden is merely shifted in the network, since a high percentage of the pre-empted calls are retried by lower-privileged sub-

scribers. An additional network impact of this feature is that users with high precedence privileges tend to call each other at that high precedence level, thus encountering the same number of station busies (non-pre-emptable) as calls at lower precedence.

5.2.3.2.2 Conferencing - Conferencing tends to have beneficial effect on network performance, since discussion with many subscribers can be held over a period of time less than the sum of that for individual discussions. The increased signaling time for sequential calling of conferees is usually small in relation to the voice holding time on trunks.

5.2.3.2.2.1 Progressive Conference - Privileged subscribers or the operator should have the capability of progressively calling subscribers into a conference call. If each call made to bring a new conferee uses another trunk, then this can have a significant impact on the DAX network loading. As with a preprogrammed conference, this feature is usually operated at a high precedence level in order to increase the probability of reaching all conferees using pre-emption to accomplish this, if necessary). The time needed to dial up and contact subscribers is greater than that of the preprogrammed conference, where it is all done automatically, and trunks and links are thus busied longer in the network.

5.2.3.2.2.2 Preprogrammed Conference - A preprogrammed conference, initiated by removing the handset or by dialing a special code, is automatically completed by the switch to a number of predesignated subscribers. Since the demand for trunks and processing loading is simultaneous rather than random, a heavy instantaneous signalling burden may be placed on the DAX network and on the DAX itself during conference setup. The originating subscriber has available his highest allowed precedence to provide protection for the conference circuit and all conferees. This level of precedence should be relatively high since the value of a preprogrammed conference would be greatly reduced if the probability of reaching all conferees at once were low.

5.2.3.2.2.3 Broadcast (and Meet-me) Conferences - The network impact of these services is similar to those progressive and preprogrammed conferences. The broadcast feature is different from the other conference features only in that it may have a much greater number of conferees.

5.2.3.2.3 Call Transfer/Call Forwarding - When DAX subscribers are temporarily located at another station, incoming calls can be forwarded to a selected station within a local area, or by an outgoing trunk to a distant location. The latter situation implies a network impact. This feature has the beneficial effect of allowing the user to complete the call either to the desired party or to someone who can either take messages or provide the pertinent information. This diminishes the number of call repeat attempts in the DAX network.

5.2.3.2.4 Compressed and Abbreviated Dialing - The main impact of this feature is in the reduced holding time for registers at the DAX itself. This feature is usually provided for local nodal subscribers only, and thus impacts at the switch, not the network.

5.2.3.2.5 Secure Call Mode Conversion - The requirement for interoperability with a wide variety of digitized voice instruments will impact the DAX itself (as discussed in Section 5.1) and the DAX network. All secure calls requiring network access must be routed over approved or approrpiately encrypted trunks. The requirements for security will impact the network throughput as the requirement for security will be conveyed from node to node via CCIS traffic.

5.2.3.2.6 Automatic Line/Trunk Grouping and Hunting - When a called station is not idle or pre-emptable, the call can be routed to an alternate station. Ths hunting may be performed in a sequential, predetermined or random fashion either by equipments or by directory number. This feature decreases the number of call repeat attempts since there is a high likelihood of being connected to someone who can either take a message or provide the desired information.

5.2.3.2.7 Direct Access Service - The main impact of this feature is the reduced demand on signalling registers at the DAX itself. The network impact is the same as for ordinary call placement.

5.2.3.2.8 Attendant Recall/Operator Functions - The attendant operator can be used to provide a wide variety of subscriber requests and services. These include providing general call assistance, trouble reporting, information, or last resort routing of the switching equipment. There are two types of operator calls; local and remote. Only the remote operator calls will affect the network.

Used effectively, attendant recall with attendant/operator service can reduce the number of wrong calls in the network, thereby increasing the grade-of-service. Steps should be taken, however, to ensure that the increased volume of traffic in the network, due to operator service, is minimal.

Calls to the remote operator can cause an increase in traffic on the links as well as at the switch. There is also the possibility of tying up trunk links for some time while the call is being assisted. This occurs since the use of an operator ordinarily increases the time necessary to process a call.

5.2.3.2.9 Data Service - The DAX will provide real-time data switching for certain bulk data subscribers. Assigning a Class I time slot for data subscribers with large amounts of bulk data will increase the network throughout since individual packet overheads can be eliminated. However, the adaptive routing aspects of the network will not then be utilized.

5.2.3.2.10 Intercept - An intercept feature provides for the return of a recorded announcement, or connection to the attendant when the number dialed does not exist, is unassigned, or is marked disabled. Although it can tie up network trunk circuits briefly, this feature can often serve to keep users from repeating call attempts by providing specific information when the call is not completed. This improves the network grade-of-service.

5.2.3.3 Impact of Class II Service Features

5.2.3.3.1 Message Accountability - The DAX network concept for message accountability calls for packet accountability at the tandem DAX's and full message accountability at the originating DAX and the terminating DAX. Packet numbering and network bookkeeping will impact the network throughput. However, high throughput alone is of no value if packet/message accountability is of prime concern.

5.2.3.3.2 Message and Communication Line Security - Digital transmission facilities of the future DCS will be encrypted with equipments of the TENLEY type. There is a significant network impact of providing this feature. The main impacts are the cost of the cryptographic equipments required for transmission COMSEC and the network management required to ensure that communications links requiring security receive it without an excessive amount of key changing and digital processing/synchronization delays.

5.2.3.3.3 Precedence and Pre-emption - The main network impact of Class II precedence and pre-emption processing is that of providing Network Control/Management for tandem traffic. Providing this feature has an impact on the control/memory architecture for the backbone DAX network. As the system architecture evolves, the intent is to minimize the complexity of these pre-emptive rules in order to minimize the impact on the DAX network processing load (time burden), while providing a cost-effective network pre-emptive service.

5.2.3.3.4 Multiple Addressing - Multiple addressing will have a significant impact on the DAX network storage and processing capability. A network handling multiply-addressed messages will have to be designed differently from single source-to-single destination flow control such as currently exists in the ARPA network.

Multiple addressing will impact the network grade-of-service (especially in the case of the geometric progression of linked multiply addressed messages).

Packet switching techniques do not lend themselves to elegant solutions of the multiple message copy issue. Pure packet switches generally provide no long-term storage, so implementation of the store-and-forward automatic multiple-message generator technique is not consistent with the design objectives of pure packet switching.

The likely alternative is to provide more buffer storage at nodes -- sufficient high-speed storage for short, top-priority messages, and disk or tape - slower but cheaper - for longer messages that can tolerate more delay. A network of this type effectively combines packet switching with customary store-and-forward operation.

5.2.3.3.5 Universal Data Subscriber Interface - As discussed in Section 5.1.3.2.5 the TRI-TAC Data Adapter (DA) is a representative universal data subscriber interface. Through the use of a hardware/software processing function the DAX must interface the DA'a operating at several transmission rates, with or without forward error correction (of several different types), with or without multisampling, or with both forward error correction and multisampling combined. This data communication flexibility would impact the nodal/network message processing load and certainly would increase the network "cost".

5.2.3.3.6 Mode/Code/Speed/Format Conversion - Since the DAX will be capable of interoperability with a wide variety of data terminal/sources, the amount of the mode/code/speed/format conversion provided will certainly impact the nodal/network message processing load by increasing the number of potential connections among data networks.

5.2.3.3.7 Archival Storage, Intercept, and Retrieval and Trace - Among the many common message handling requirements in a strategic/tactical message switching environment are the capability for archival storage, message retrieval and trace, and temporary interception of all traffic for specific routing indicators. They are treated together here, since the network solution for these services is similar. Regional archive and intercept centers, along the AUTODIN line, having enough bulk storage to satisfy the storage function for all the DAX's in a given region, could be provided. Data communication protocols would then have to be developed to remotely store and access from these centers. Network "cost" would certainly be increased by providing this communication and by providing the bulk memory storage required at the regional DAX centers.

5.3 SYSTEM SERVICE FEATURES

5.3.1 Problem

System control management service features are those services not directly associated with individual subscribers but related to the overall control of the switching center and its related switching network. The concepts associated with DAX system control/management are illustrated in Figure 5-1. The problem is to assess the applicability and impact of these system/control management service features on the DAX and the DAX network.

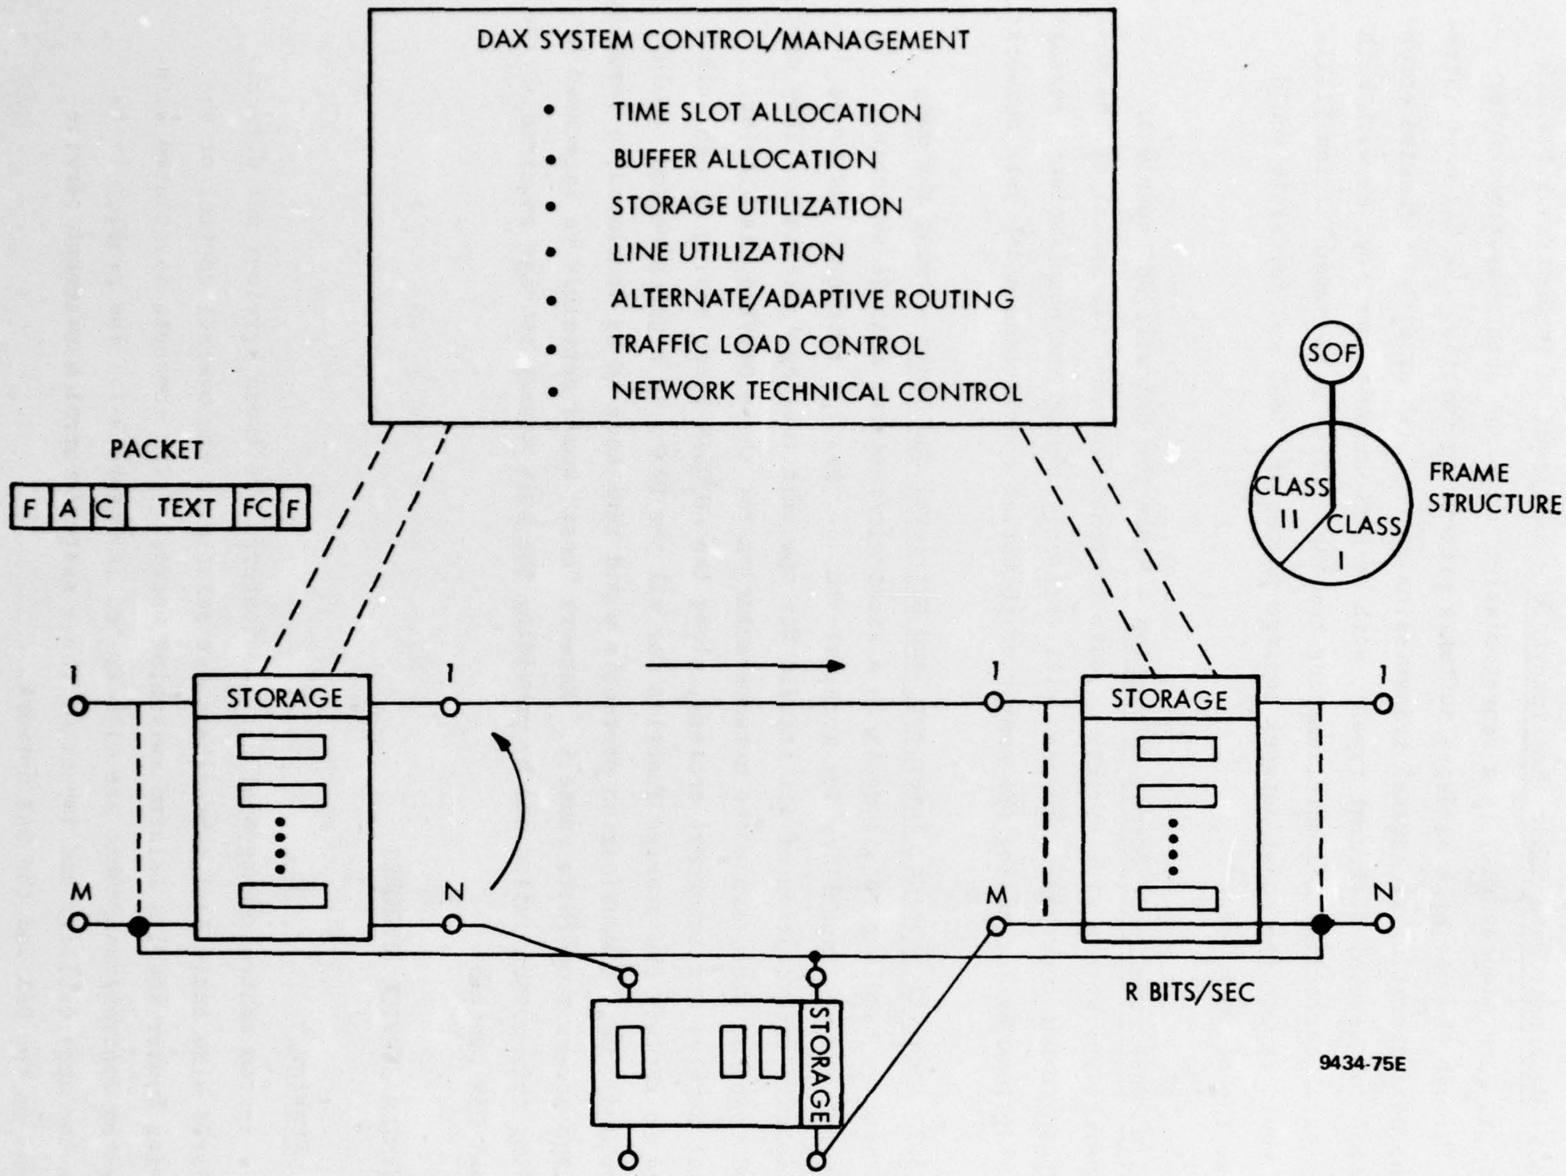

Figure 5-1. DAX System Control/Management Concepts

5.3.2 Objective

The objective of this section is to assess the desirability and feasibility and impact of providing system service features at the DAX. Three main system service features will be assessed:

a. Alternate/Adaptive Routing

b. Traffic Load Control

c. Network Technical Control Functions.

5.3.3 Progress

5.3.3.1 Alternate/Adaptive Routing

Switching network routing schemes recommended for the SENET-DAX concept can be divided into two categories:

a. Alternate Routing (Class I) - based on a fixed average input traffic matrix of the network

b. Adaptive Routing - (Class II) - based on the dynamic load of the network.

5.3.3.1.1 Class I Traffic - Alternate Routing - The routing scheme proposed for Class I traffic (Sections 3.1 and 3.2) incorporates an alternate routing capability accomplished through the use of routing tables stored at each DAX. These tables define primary and alternate routes from a given switch to every other switch in a network. A Class I call will always be routed over the primary route if possible, then the first alternate, etc. The routing table at each switch is unique to that switch and is fixed, i.e., the primary route to a given destination switch is always the primary route, the first alternate is always the first alternate, etc.

The size of these tables will determine the amount of memory required at each DAX. Also the size of the tables will determine the complexity of the search algorithm and, hence, the amount of processing time consumed by the switch processor.

5.3.3.1.2 <u>Class II Traffic - Adaptive Routing</u> - The routing scheme proposed for Class II traffic will incorporate an adaptive routing capability. Class II packets will be directed to the route which has the smallest predicted delay from the DAX in question to the destination DAX. The predicted delays to each destination switch will be contained in an adaptive routing table which is updated periodically. A technique being considered is an adaptive routing algorithm similar to that used in the ARPA network. The ARPA network operates with only two levels of precedence: priority for single-packet segments (which includes system control packets), and no priority for everything else. The ARPA algorithm will require modification to handle the multi-level DAX precedence and preemption. For example, the ARPA algorithm uses queuing delays as a factor in calculating predicted overall delays. In a network carrying data traffic of various precedence levels, the actual queuing delay would be dependent on the precedence level of the packet being routed. Further analysis of these priority queuing delays will be necessary.

Again, the size of the adaptive routing table will determine the amount of memory required at each switch. The size of the tables will determine the complexity of the search algorithm and, hence, the amount of processing time consumed by the processor. The frequency of adaptive routing table updating will also affect the DAX processing overhead.

5.3.3.2 Traffic Load Control

There are in principle two different modes of operating telecommunications switching networks: free-wheeling traffic control and complete traffic control. We say "in principle", because in practice a network is generally operated by a combination of these modes.

In a free-wheeling environment a terminal or a switching center is able to generate or forward data at any time and the receiving side has to be ready at all times to accept this data, i.e., the receiving side is able to cope with every traffic peak without rejecting any incoming traffic. The alternative mode of operating a telecommunications facility is to completely control the amount of incoming traffic. A typical example of complete traffic control is a centralized computer network, in which widely spread terminals are polled directly by the central processor. Terminals which have messages to send have to wait until they

are solicited, that is, until the receiving center signals that it is ready to accept traffic. In the event that the processor is nearing an overload situation, in terms of running short of storage or of processing capacity, it is able to reduce the polling rate, thereby slowing down the receipt of messages, or it can stop polling completely.

The traffic load control procedure for the DAX network will of course lie someplace between complete traffic control and free-wheeling traffic control. By monitoring a switch's send, receive and task queues, buffers, storage and line utilization a traffic load control procedure can be implemented. Note that Class II (packet/message) traffic will be throttled (delayed) and Class I (voice) traffic will be blocked by denying access to the switch either at the terminal or the trunk/time slot allocation level. Further analysis will be required to determine the optimum traffic load control strategy.

5.3.3.3 Network Technical Control Functions

The technical control functions performed by the DAX will probably be of the communication equipment support type. The normal interface of this function with transmission equipment is through a communications nodal control function; however, it should have the capability to interface switched trunks and groups directly with transmission equipment. When the DAX is in an emergency bypass condition, the equipment support function should probably interface up to 10 percent of the circuits directly to the transmission equipment.

The DAX equipment support function will be responsible for patching, executing reconfiguration, restoration, testing and status reporting of all loops and switched trunks and groups. Line conditioning should also be provided by this function.

The main impact of these technical control functions is the memory storage required for diagnostic/status reporting software, redundant hardware for rapid restoration, and the transmission overhead required to communicate the status of the DAX to the next level in the technical control hierarchy.

5.3.4 Conclusions

It appears that the use of system control and service features discussed in this section will impact processor loading and memory requirements at each switch, and the amount of control overhead sent through the network. It appears that they will also affect time slot allocation, buffer allocation, use of control storage, and the kind and magnitude of traffic measurements, among others.

Other system control features requiring investigation in the future are remote data base interrogation and interswitch master frame synchronization. Remote data base interrogation and change from remote Syscon centers is obviously a desirable system feature. The ability of the remote centers to synchronize master frames among the DAX switches, if feasible, may decrease network call completion times and data transfer delays.

System control is obviously a function of the switching network structure and complexity, as well as of the routing scheme, anticipated magnitude of data changes and traffic reporting, required system availability, and the like. A reasonable system control structure, hierarchical if that seems indicated, must be hypothesized in order to perform further analysis. Sophisticated computer processing and displays at the Syscon centers must also be considered.

Further development of system and technical control in a SENET-DAX network system also requires that the nature, format, and periodicity of the basic data base information and changes that must be provided to and received from system control centers must be determined. In addition, the accumulation of traffic data, its preprocessing in a background mode, and its distribution to system control centers is also a consideration. This information exchange will impact the frame contents in the basic switching concept, and the software required for associated data processing. The effect of the alternate/adaptive routing technique on the switching concept, including software impact, must also be studied further.

A reasonable digital interface between a DAX and a Syscon control center is a necessity. Initially, Syscon data will be assumed to be routed among switches over the common-user trunks as packetized data, with dedicated Syscon channels between a Syscon center and its associated DAX.

Some degree of technical control, interrelated with system control and integrated switch control itself, is anticipated to be required. This technical control will incorporate measurement, although not necessarily feedback control, of transmission quality parameters such as noise, signal levels, delays, etc. Distribution of the data obtained will be made in the same manner as system control data.